

CompTIA A+® Certification
ALL-IN-ONE DESK REFERENCE
FOR
DUMMIES®

by Glen E. Clarke and Ed Tetz



Wiley Publishing, Inc.

CompTIA A+® Certification All-in-One Desk Reference For Dummies®

Published by
Wiley Publishing, Inc.
111 River Street
Hoboken, NJ 07030-5774

www.wiley.com

Copyright © 2007 by Wiley Publishing, Inc., Indianapolis, Indiana

Published by Wiley Publishing, Inc., Indianapolis, Indiana

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 either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600. Requests to the Publisher for permission should be addressed to the Legal Department, Wiley Publishing, Inc., 10475 Crosspoint Blvd., Indianapolis, IN 46256, (317) 572-3447, fax (317) 572-4355, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, the Wiley Publishing logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies Daily, The Fun and Easy Way, Dummies.com, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. CompTIA A+ is a registered trademark of the Computing Technology Industry Association. The CompTIA Authorized Quality Curriculum logo is a proprietary trademark of CompTIA. All rights reserved. All other trademarks are the property of their respective owners. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, please contact our Customer Care Department within the U.S. at 800-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002.

For technical support, please visit www.wiley.com/techsupport.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Library of Congress Control Number: 2006936832

ISBN: 978-0-471-74811-3

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

10/RQ/QR/QX/IN



About the Author

Glen E. Clarke (MCSE/MCSD/MCDBA/MCT/CIW SA/Security+/Network+/A+) is an independent trainer and consultant, focusing on network security and security auditing services. Glen spends most of his time delivering certified courses on A+, Network+, Windows Server 2003, SQL Server, Exchange Server, Visual Basic .NET, and ASP.NET. Glen also teaches a number of security related courses covering topics such as vulnerability testing, firewall design, and packet analysis.

Glen is an experienced author and technical editor who has worked on nine certification books. Glen designed and coauthored the award nominated *A+ Certification Bible* and has worked on certification titles involving topics such as Windows 2000/2003 certification, CIW certification, Network+ certification, and Security+ certification.

When he's not working, Glen loves to spend quality time with his wife, Tanya, and their three children, Sara, Brendon, and Ashlyn. He is an active member of the martial arts community, where he currently holds his first-degree black belt in Tae Kwon Do. You can visit Glen online at www.glenclarke.com, or contact him at glenclarke@accesswave.ca.

Ed Tetz graduated in 1990 from Saint Lawrence College in Cornwall, Ontario with a degree in Business Administration. He spent a short time in computer sales, which eventually led to a computer support position. After several years of providing system and LAN support to small and large organizations, in 1994 he added training to his repertoire. He holds certifications for A+, ITIL Foundations Certificate in IT Service Management (IT Infrastructure Library), Microsoft Certified Trainer (MCT), Microsoft Certified Systems Engineer (MCSE), Microsoft Certified Database Administrator (MCDBA), and Chauncey Group's Certified Technical Trainer (CTT). Since 2002, he has been a full-time consultant for a value added reseller in Halifax, Nova Scotia. Over his years of work experience, he has supported Apple Macintosh, IBM OS/2, Linux, Novell NetWare, and all Microsoft operating systems from MS-DOS to Windows Vista, as well as hardware from most of the major vendors. He welcomes comments from his readers and can be contacted at ed_tetz@hotmail.com, if you are not trapped by the junk mail filters, or info@edtetz.net.

Dedication

To my beautiful wife, Tanya, who has made all my dreams come true.
I cherish every moment we spend together.

— Glen E. Clarke

I would like to dedicate this book, with love, to my wife, Sharon, and my children, Emily and Mackenzie. They have put up with a lot during the writing of this book, especially the loss of my time, which is now gone forever. If I owe anyone my gratitude for having this book written, it is them.

— Ed Tetz

Author's Acknowledgments

There are so many hard working people that are needed to write a book and these people never get the credit they deserve — so I will attempt to show my appreciation. I want to thank the people at Wiley for their hard work and continued support. A special thank you to acquisitions editor, Katie Feltman, for asking Ed and I to do this book and for her patience as I worked through the chapters. I also want to thank project editor Blair Pottenger for his patience and quick responses to any questions I fired his way! Much appreciation goes to copy editor Andy Hollandbeck for an outstanding job on the editing of the chapters — you have made a *huge* difference in this project — thank you! I also want to thank Dan DiNicolo for his feedback as the technical editor.

Over the years I have developed friendships through my work, and one of my favorite people to work with is Ed Tetz. I want to thank him for the long hours he has dedicated to this book — you have done an amazing job again! I also want to thank Dan Lachance, who is always around for a little bit of geek talk! Thanks to Riley, Haley, and Sara for helping me take some of the photos for the book.

A special thank you goes to my three children, Sara, Brendon, and Ashlyn for giving me the time to sit down and participate in kids play with them — it makes me realize what is important in life. A needed thank you goes to my wife, Tanya, for all her support — without it I would be a different person. Thank you for all that you do!

— Glen E. Clarke

I would like to thank Katie Feltman, our acquisitions editor, for inviting me to take on this project, and to both Katie and Blair Pottenger, our project editor, for keeping this book on track, which I am sure was an effort for them, and thanks for that. I would like to thank Andy Hollandbeck, our copy editor, for his tireless work of ensuring that what I wrote made sense, and for making my words sound better than when I originally wrote them.

In addition to them, I would like to thank the rest of the staff at Wiley Publishing who worked behind the scenes taking care of many of the details that are required to get this book into print, and I am sure that due to their hard work, this will be a successful book.

— Ed Tetz

Publisher's Acknowledgments

We're proud of this book; please send us your comments through our online registration form located at www.dummies.com/register/.

Some of the people who helped bring this book to market include the following:

Acquisitions, Editorial, and Media Development

Project Editor: Blair J. Pottenger
Acquisitions Editor: Katie Feltman
Copy Editor: Andy Hollandbeck
Technical Editor: Dan DiNicolò
Editorial Manager: Kevin Kirschner
Media Development Specialists: Angela Denny,
Kate Jenkins, Steven Kudirka, Kit Malone
Media Development Coordinator:
Laura Atkinson
Media Project Supervisor: Laura Moss
Media Development Manager:
Laura VanWinkle
Editorial Assistant: Amanda Foxworth
Sr. Editorial Assistant: Cherie Case
Cartoons: Rich Tennant (www.the5thwave.com)

Composition Services

Project Coordinator: Erin Smith
Layout and Graphics: Claudia Bell,
Lavonne Cook, Denny Hager,
Joyce Haughey, Stephanie D. Jumper,
Barbara Moore, Barry Offringa
Special Art:
Proofreaders: John Greenough,
Christy Pingleton
Indexer: Richard Shrout
Anniversary Logo Design: Richard Pacifico

Publishing and Editorial for Technology Dummies

Richard Swadley, Vice President and Executive Group Publisher
Andy Cummings, Vice President and Publisher
Mary Bednarek, Executive Acquisitions Director
Mary C. Corder, Editorial Director

Publishing for Consumer Dummies

Diane Graves Steele, Vice President and Publisher
Joyce Pepple, Acquisitions Director

Composition Services

Gerry Fahey, Vice President of Production Services
Debbie Stailey, Director of Composition Services

The logo of the CompTIA Authorized Quality Curriculum (CAQC) program and the status of this or other training materials as “Authorized” under the CompTIA Authorized Quality Curriculum program signifies that, in CompTIA’s opinion, such training material covers the content of CompTIA’s related certification exam.

The contents of this training material were created for the CompTIA A+ Certification exam covering CompTIA certification objectives that were current as of 2006.

CompTIA has not reviewed or approved the accuracy of the contents of this training material and specifically disclaims any warranties of merchantability or fitness for a particular purpose. CompTIA makes no guarantee concerning the success of persons using any such “Authorized” or other training material in order to prepare for any CompTIA certification exam.

How to become CompTIA certified:

This training material can help you prepare for and pass a related CompTIA certification exam or exams. In order to achieve CompTIA certification, you must register for and pass a CompTIA certification exam or exams.

In order to become CompTIA certified, you must:

- 1.** Select a certification exam provider. For more information please visit http://www.comptia.org/certification/general_information/exam_locations.aspx.
- 2.** Register for and schedule a time to take the CompTIA certification exam(s) at a convenient location.
- 3.** Read and sign the Candidate Agreement, which will be presented at the time of the exam(s). The text of the Candidate Agreement can be found at http://www.comptia.org/certification/general_information/candidate_agreement.aspx.
- 4.** Take and pass the CompTIA certification exam(s).

For more information about CompTIA’s certifications, such as its industry acceptance, benefits, or program news, please visit www.comptia.org/certification.

CompTIA is a not-for-profit information technology (IT) trade association. CompTIA’s certifications are designed by subject matter experts from across the IT industry. Each CompTIA certification is vendor-neutral, covers multiple technologies, and requires demonstration of skills and knowledge widely sought after by the IT industry.

To contact CompTIA with any questions or comments, please call 1-630-678-8300 or email questions@comptia.org.

Table of Contents

<i>Introduction</i>	1
About This Book.....	1
Conventions Used in This Book	1
Foolish Assumptions	2
How This Book Is Organized.....	3
Book I: Setting the A+ Groundwork.....	3
Book II: Inside the Box	3
Book III: Outside the Box.....	3
Book IV: Maintenance and Troubleshooting.....	3
Book V: Operating System Basics.....	3
Book VI: Managing the Operating System	4
Book VII: Recovering Systems	4
Book VIII: Networking	4
Book IX: Securing Systems	4
Appendixes.....	4
Icons Used in This Book.....	4
Where to Go from Here.....	5
<i>Book I: Setting the A+ Groundwork</i>	7
Chapter 1: The New A+ Exams	9
CompTIA A+ Certification and Why You Need It.....	9
Checking Out the Exams and Their Objectives.....	10
The CompTIA A+ Essentials Exam	11
CompTIA A+ IT Technician Exam	12
The CompTIA A+ Remote Support Technician Exam	12
The CompTIA A+ Depot Technician Exam	13
Using This Book to Prepare for the Exams	13
Making Arrangements to Take the Exams.....	14
The Day the Earth Stood Still: Exam Day	14
Arriving at the exam location	14
Taking the exam.....	15
How does CompTIA set the pass level?.....	17
Chapter 2: A+ Soft Skills	19
Using Troubleshooting Procedures and Good Practices.....	19
Identifying the problem	21
Analyzing the problem and potential causes	22
The basic troubleshooting process	22
Evaluate results	25

Document findings, activities, and outcomes.....	25
Documentation resources	26
Professionalism and Communication.....	27
Good communication skills.....	27
Professional behavior	28
Getting an A+	31
Chapter 3: Protecting Yourself and Your PC	35
Gathering Tools of the Trade.....	35
Hardware tools	35
Software tools	38
Avoiding the Dreaded ESD	38
ESD means electrostatic discharge (not extrasensory deprivation)	38
Preventing ESD.....	39
Watching Out for EMI.....	40
Recognizing High-Voltage Equipment.....	40
Power supply unit (PSU)	41
Monitor	41
Disposing of Components	42
Batteries	42
CRTs	43
Toner kits and cartridges	43
Material Safety Data Sheet (MSDS)	43
Exploring Ergonomics	44
Getting an A+	45
Chapter 4: An Overview of System Components	49
What Is a Computer?.....	49
Looking Inside the Box	50
Processor/CPU.....	50
Storage devices.....	51
Memory.....	52
System boards	53
Power supply	53
Adapter cards	53
Cooling system and fans.....	54
Firmware and chipsets	54
BIOS.....	55
CMOS.....	55
Checking Outside the Box.....	56
Casing and form factors	56
Input and output devices	56
Getting an A+	59

***Book 11: Inside the Box* 63**

Chapter 1: Knowing Your Motherboard 65

Finding Out What’s on a Motherboard	66
Processor	66
SIMM/DIMM sockets	68
Cache memory	69
Motherboard chipset	70
BIOS chip	72
Battery	72
Expansion slots	73
Ports and connectors	74
Power connectors	82
Drive connectors	85
Jumpers and DIP switches	88
Identifying the Types of Motherboards	89
Full AT	90
Baby AT	91
LPX/NLX	93
ATX	93
MicroATX and FlexATX	95
Understanding Bus Architectures	96
ISA	97
MCA	97
EISA	98
VESA	100
PCI	101
PCMCIA	102
AGP	103
PCI-X	103
PCI Express	104
AMR and CNR	105
Performance Considerations	106
Getting an A+	107

Chapter 2: Picking Your Processor 115

Understanding Processor Terminology	115
Processor speed	116
Data bus	116
Address bus	117
Registers	118
Cache memory	118
Math co-processor	119
Real-mode versus protected-mode	120
MMX	120
Hyperthreading	121
Dual core processors	121
Throttling	121

Overclocking	122
VRM	122
Chip packaging	122
Identifying Socket Types	125
Looking at Popular Intel Processors.....	127
Pentium.....	127
Pentium Pro.....	130
Pentium II.....	131
Celeron.....	131
Pentium III	132
Xeon	133
Pentium 4.....	134
Itanium and Itanium II.....	134
Pentium “M”	135
Don’t Forget Non-Intel Chips	136
K6.....	136
K6-2.....	136
K6-III	136
Athlon	137
Athlon XP.....	137
Duron	137
Opteron.....	137
Installing a Processor	138
Will it fit in the socket?	138
CPU voltage and transistor integration	138
Performing the installation	138
Keeping a Processor Cool	141
Heat sinks and CPU fans	142
Installing a heat sink and fan	142
Increasing Performance	143
Getting an A+	144

Chapter 3: What to Remember about Memory151

Understanding the Types of Memory.....	151
Remembering the purpose of memory.....	151
Read-Only Memory (ROM).....	152
Random Access Memory (RAM)	153
Identifying the Types of DRAM.....	156
Standard DRAM.....	156
Fast page mode.....	157
Extended data output	158
Burst Extended Data Output.....	158
Synchronous DRAM	158
Rambus DRAM	159
DDR.....	159
DDR2.....	159
How Would You Like Your Chips Packaged?.....	160
SIMMs.....	160
DIMMs	161

SODIMM	162
MicroDIMM.....	162
Understanding Error-Checking Memory	163
Parity versus non-parity.....	163
ECC memory	164
Working with Cache Memory	164
L1 cache.....	165
L2 cache.....	165
L3 cache.....	166
Installing or Upgrading Memory	166
Type of memory	166
Speed.....	166
Connectors	167
Parity versus non-parity.....	167
Installing memory on desktop PCs	168
Installing memory on laptop systems	169
Getting an A+	171

Chapter 4: Telling Your BIOS from Your CMOS179

The BIOS and Its Purpose	179
Upgrading the System BIOS	181
Performing the BIOS upgrade	181
Potential issues with BIOS upgrading.....	182
Understanding CMOS	183
Viewing Basic CMOS Settings	185
Hard drive.....	185
Floppy disk drive.....	187
Memory.....	187
Parallel ports.....	188
Serial ports	190
Date and time	190
Boot sequence	191
Passwords	192
Plug and Play BIOS	193
Viewing Advanced CMOS Settings	194
Globally Unique Identifier (GUID)	194
BIOS date and revision number.....	194
Universal Serial Bus (USB)	195
Built-in network adapter.....	196
Virus protection.....	196
On-board cache	197
Reserve resources	198
Getting an A+	199

Chapter 5: Working with Storage207

Understanding Hard Drive Terminology	207
Disk geometry.....	207
Read/write process	210
Performance.....	211

Master Boot Record	212
LBA and ECHS	212
Discovering IDE Devices.....	213
IDE overview	213
Installing IDE devices	217
Learning How SCSI Works	222
SCSI overview.....	222
Installing SCSI devices	230
The Serial ATA Hype!.....	232
Using Removable Storage	233
Floppy disks	233
CD-ROM/CDRW	234
DVD/DVDRW	235
Flash drives	235
USB external drives	235
Zip drives.....	236
Flash cards	236
Tape.....	236
Understanding File Systems	237
The FAT file system	237
The FAT32 file system	239
NTFS	240
NTFS 5.0	240
HPFS	241
Managing Partitions and Volumes	241
Primary partition.....	242
Extended partition	242
Creating partitions and volumes in Windows 2000/XP/2003	244
Formatting partitions and volumes	250
Securing Data with RAID	252
Mirroring/duplexing (RAID level 1)	253
RAID 5 Volume (RAID level 5)	257
Understanding Management Tools	259
Defragmentation utility.....	260
Check Disk utility.....	261
Disk Cleanup	262
Getting an A+	264
Chapter 6: Working with Power	271
Knowing the Basics of Power Terminology	271
Volt	271
Ohm.....	272
Amp	272
Watt	272
Identifying the Purpose of Power Supply Units	272
Identifying Power Supplies	274
Power Connectors.....	276
The ATX main power connector	276
The floppy drive connector	277

The peripheral connector	278
The 12-volt power connector.....	278
The Serial ATA connector	279
Power Supply Form Factors.....	279
Using AC Adapters	280
Working with UPS and Suppressors.....	282
Getting an A+	284

***Book III: Outside the Box* 289**

Chapter 1: Ports, Cables, and Connectors 291

Identifying Common Computer Ports.....	291
Serial and parallel ports	292
Universal Serial Bus (USB)	293
FireWire (IEEE-1394).....	293
Keyboard	294
Monitor	294
Comparing Cable Types	294
Ribbon.....	295
Twisted pair	295
Thick and thin coax.....	297
Fiber	297
Cable Orientation.....	298
Connector Types	299
IBM Type 1 Connector	299
DB-9	300
DB-15	300
DB-25	300
Centronics 36 and 50.....	302
RJ-11	303
RJ-45	303
BNC.....	303
PS/2 or Mini-DIN 6.....	304
Universal Serial Bus (USB) connectors	304
IEEE-1394 (FireWire) connectors	305
Standard External Cables	306
Parallel Cable	306
Serial Cable.....	307
Null Modem Cable	307
Viewing Cable Adapters	308
Barrel connectors.....	308
Gender changers	308
Null modem.....	309
Multimedia Connectors	310
Getting an A+	311

Chapter 2: Installing and Configuring Input Devices	315
Minding Your Keys and Qs.....	315
Identifying keyboard types.....	317
Installing a keyboard.....	318
Configuring keyboards.....	318
Catching the Mouse	321
Types of mice.....	321
Installing a mouse.....	322
Communicating with Modems and Network Adapters	323
Working with modems	323
Working with network adapters	326
Other Input Devices	329
Getting an A+	330
Chapter 3: Installing and Configuring Output Devices	335
Understanding Video Adapters	335
Looking at the video standards	336
Video board features.....	337
Installing a video card.....	339
Using Your Monitor	339
Types of displays.....	340
Configuring your display settings	341
Configuring multi-display support	342
The Sound of Computers	343
Types of sound cards.....	343
Installing a sound card.....	344
Configuring a sound card	345
Other Output Devices.....	346
Getting an A+	347
Chapter 4: Examining System Resources	351
Understanding System Resources	352
I/O addresses	352
Interrupt ReQuest (IRQ)	355
Direct Memory Access (DMA)	360
Memory addresses.....	363
Working with System Resources	363
Identifying resource conflicts	364
Changing system resources	365
Getting an A+	371
Chapter 5: Managing Printers	379
Learning about Paper Feeder Mechanisms	379
Continuous form feeders.....	380
Friction feeders.....	380

Understanding Types of Printers	380
Laser printers	380
Inkjet printers	385
Dot matrix printers	386
Thermal printers	387
Viewing Types of Printer Connections and Configurations.....	387
Parallel connections.....	387
Serial cable connections.....	388
Network cable connections	388
Universal Serial Bus (USB)	389
Infrared	390
FireWire and SCSI	390
Upgrading Printer Hardware.....	391
Upgrading memory	391
Adding a disk drive	391
Upgrading the firmware.....	391
Installing a Printer in Windows	392
Installing a printer	392
Configuring a printer.....	394
Connecting to the shared printer.....	396
Troubleshooting Printer Problems	397
Check the simple stuff first	397
Paper jams.....	398
Garbled or corrupted output.....	398
Spots or smudging on the printout.....	398
Slow printing.....	398
Spooling service problems.....	399
Dot matrix problems	400
Inkjet problems.....	401
Laser printer problems.....	401
Understanding Safety and Preventative Maintenance	402
Safety precautions.....	402
Preventative maintenance.....	402
Getting an A+	403

Chapter 6: Working with Multimedia Devices 407

Understanding Scanners	407
Scanning process.....	409
Connecting a scanner	409
Scanning an image.....	409
Using Digital and Web Cameras.....	411
Digital cameras	411
Uploading pictures to the computer.....	413
Web and digital video cameras.....	413
Looking at Other Multimedia Devices	416
Microphones	416
MIDI	416
Getting an A+	417

Chapter 7: Dealing with Portable Computers	421
Identifying Portable Computer Components.....	421
Looking at Laptop Batteries	423
Different types of batteries	423
Handling batteries	424
Maximizing battery performance.....	424
Changing batteries	425
Power management features	426
Understanding AC Adapters	426
AC adapter problems	427
AC adapter troubleshooting and repair	427
Learning about LCD Panels.....	427
Handling LCD panels.....	428
Connecting an LCD panel to a computer.....	429
Understanding Laptop Input Devices.....	429
Laptop keyboard	429
Touch pad and the rubber mouse ball	430
Laptop Communication Components.....	432
Network card and modem.....	432
Wireless network card	432
Other communications ports	432
Port Replicators and Docking Stations	433
Expanding on a Laptop.....	434
Adding a PC card	435
Removing a PC Card.....	435
Upgrading Your Laptop	436
Upgrading memory	436
Adding hard disk space	437
Understanding Handheld Devices.....	438
Getting an A+	439

Book IV: Maintenance and Troubleshooting**445**

Chapter 1: Performing Preventative Maintenance	447
Preventative Maintenance 101	447
Using Preventative Maintenance Tools	449
Materials and equipment	449
Software utilities.....	450
Maintaining Environmental Controls.....	450
Ventilation and airflow.....	450
Humidity and liquids.....	454
ESD	454
Dirt and dust	455
EMI.....	455
Power, UPS, and suppressors	456

Completing Maintenance Tasks	457
Case and components.....	457
Power supplies	458
Monitors	458
Keyboards and mice	459
Drives	459
Storage media	461
Laptops	461
Printers and scanners.....	462
Getting an A+	463

Chapter 2: Troubleshooting Procedures and Guidelines 467

Identifying Troubleshooting Tools.....	467
Hardware tools	467
Diagnostic software.....	473
The Art of Troubleshooting	474
Troubleshooting basics	475
Physical environment	475
Audio and visual troubleshooting.....	475
CMOS and BIOS.....	476
Motherboard.....	477
Processor/memory.....	477
Floppy drive	478
Hard drives.....	478
CD/DVD-ROM	480
Keyboard and mouse	481
Sound card/audio	482
Monitor/video	482
Modem	484
Serial and parallel ports	484
USB	485
Power supplies and batteries	487
Laptops	489
Getting an A+	493

Book V: Operating System Basics 497

Chapter 1: Comparison of Major Operating Systems 499

What Is an Operating System?.....	499
Looking at the Windows OS	500
Comparing versions of Windows.....	500
Typical system requirements	501
Identifying GUI elements	502
Welcome to the Linux World.....	503
Comparing versions of Linux.....	503
Typical system requirements	505
Identifying GUI elements	506

Understanding the Macintosh OS	507
Comparing versions of Mac OS	507
Typical system requirements	508
Identifying GUI elements	508
Getting an A+	509

Chapter 2: Operating System Functions 513

Identifying Major Operating System Functions.....	513
Understanding memory management	514
Checking the OS version	514
Understanding Major Operating System Components	517
Getting into the architecture	517
Paging your memory.....	518
Choosing your file systems.....	520
Registry.....	522
Navigating Your Computer.....	523
My Computer	524
Windows Explorer	525
My Network Places.....	526
Command line.....	527
Using Tools and Configuration Utilities.....	527
Control panel applets.....	527
Microsoft Management Console.....	529
Taskbar and Notification Area.....	530
Start menu.....	530
Remote Desktop Connection	530
Remote Assistance.....	531
Getting an A+	533

Chapter 3: Installing and Upgrading Windows 537

Understanding Installation Methods.....	537
Boot floppy.....	537
Bootable CD	538
Network installation.....	539
Drive imaging	539
Preparing for Installation	540
Hardware requirements.....	540
Computer name	541
Workgroup/domain	541
Partitions.....	542
File systems.....	543
Other preparation steps.....	543
Performing Attended Installations of Windows	544
Performing a Windows 2000 attended installation	545
Performing a Windows XP attended installation.....	556
Performing Unattended Installations of Windows	561
Verifying and Troubleshooting the Installation.....	564

Upgrading Windows.....	565
Preparing to upgrade.....	565
Upgrading to Windows 2000 Professional.....	566
Upgrading to Windows XP.....	567
Installing Additional Windows Components	568
Understanding How to Dual Boot Windows	570
Updating Installation Files	571
Restoring User Data Files	572
Saving a computer's state	572
Restoring a computer's state.....	574
Getting an A+	574

Chapter 4: Managing Files and Directories579

Identifying File-Naming Conventions.....	579
Long and short filenames	581
Creating file associations	583
Understanding file extensions	589
Understanding File Attributes	594
How attributes are stored	595
The basic attributes.....	596
Windows 2000 and Windows XP extended attributes	597
Setting basic attributes.....	603
Getting an A+	607

Chapter 5: Command Prompt Procedures613

Using command.com and cmd.exe.....	613
Managing Directories.....	615
dir	615
mkdir	619
chdir	619
rmdir	620
Copying and Moving Files	620
copy.....	620
xcopy.....	621
move.....	621
del.....	621
ren	622
Making Comparisons	622
diskcopy.com and diskcomp.com	622
fc.exe	622
Performing Diagnostics and Tuning Performance.....	623
defrag.exe	623
chkdsk.exe.....	624
mem.exe.....	624
Working with the Rest	625
attrib.exe	625
diskpart.exe.....	625
format.com	626

sys.com	627
ver	627
help	627
Working with the edit Command	628
Opening files	628
Saving files	629
Searching and replacing	630
Closing the editor	630
Batch File Basics	631
Starting your file	631
Getting your directions	632
Working with parameters or starting arguments	632
Making batch files make decisions	633
Looping	636
Locating a command	637
Putting it together	638
Getting an A+	638

Chapter 6: Working with System Files and the Boot Process645

Power-On Self-Test (POST) Process	646
Standard Boot Process for Windows XP	646
ntldr	647
ntbootdd.sys	648
boot.ini	648
ntdetect.com	650
ntoskrnl.exe	650
The device load process	651
The service load process	652
Loading the shell	652
Understanding the Boot Process for MS-DOS and Windows 9x	653
io.sys	653
msdos.sys	654
config.sys	656
command.com	658
autoexec.bat	659
win.com and vmm32	660
Loading the shell	662
Managing Memory	662
Conventional memory	662
Expanded memory	663
Extended memory	663
High memory	664
Upper memory	664
Virtual memory	664
himem.sys	667
emm386.exe	667
Examining Other Boot Process Files	669
smartdrv.exe	669
system.ini	671

win.ini.....	672
sysedit.exe.....	673
Getting an A+	674

Book VI: Managing the Operating System679

Chapter 1: Loading and Configuring Device Drivers681

Loading and Removing Device Drivers	681
Changing Device Settings	687
Updating Drivers	688
Signed and Unsigned Drivers.....	689
Working with Plug and Play	690
Permissions to Install Drivers	693
Verifying Driver Installation	694
Device Driver Failure	694
Getting an A+	696

Chapter 2: Working with Applications699

Installing and Removing Applications	699
Installing an application	700
Removing an application.....	703
Getting the Most out of Multitasking.....	704
Cooperative multitasking	705
Preemptive multitasking.....	705
Running 32-Bit Windows Applications	706
Benefiting from 32-bit applications	706
Executing in the Windows 9x environment.....	707
Executing in the Windows 2000 and Windows XP environments.....	708
Running 16-Bit Windows Applications	710
Executing in the Windows 9x environment.....	711
Executing in the Windows 2000 and Windows XP environments.....	712
Encountering incompatibilities	715
Running MS-DOS-Based Applications	716
Program Information File (PIF) settings	717
Understanding incompatibilities	727
Windows XP Compatibility Modes.....	728
Application Install, Start, and Load Errors	729
Getting an A+	730

Chapter 3: Optimizing the Windows Environment735

Identifying Areas of System Bottlenecks.....	735
Using Monitoring Tools	736
Task Manager	736
Performance.....	741

Optimization Best Practices	746
Virtual memory.....	746
Hard drive.....	748
Printers	748
Scanners	749
Temporary files.....	749
Services	750
Startup	754
Getting an A+	755

Chapter 4: Using Windows-Based Troubleshooting Utilities759

Understanding Operating System Utilities for Disk Maintenance	759
fdisk.exe and Disk Management	759
defrag.exe and dfrg.msc	762
chkdsk.exe (Check Disk)	764
Configuration Utilities	766
System Configuration Utility	766
regedit.exe	767
regedt32.exe	772
Device Manager	774
Computer Management	774
Getting a Handle on Diagnostic Utilities	776
The System Information Tool.....	776
Event Viewer	777
Network Diagnostics	780
Dr. Watson	780
System File Checker	781
Miscellaneous Utilities	782
Task Scheduler.....	782
Windows Script Host.....	783
edit.com	784
expand.exe	784
Getting an A+	785

Book VII: Recovering Systems **791**

Chapter 1: Managing Error Codes and Startup Messages793

MS-DOS and Windows 9x Boot Errors	793
Error in config.sys line XX.....	794
Bad or missing command.com	794
No operating system found.....	795
Windows XP Boot Files and Boot Errors.....	795
SCSI issues	796
No operating system found.....	796
Missing boot files	796

Device related errors	799
A device referenced in system.ini, win.ini, or Registry is not found	800
Registry corruption.....	800
Safe Mode	801
Other Errors.....	801
Paging file or Swap file errors	801
Failure to start GUI.....	802
Error Diagnostic Tools.....	803
Dr. Watson	803
Windows Error Reporting.....	803
Event Viewer	805
Getting an A+	806

Chapter 2: Managing Common Problems 811

Solving Windows-Specific Printing Problems	811
Dealing with a stalled print spooler	812
Incorrect/incompatible driver for printing	813
Incorrect parameters	814
Solving Boot Errors and Errors Requiring Restarting	815
Auto-restart errors	815
Blue-screen errors	817
System lockup.....	819
Resolving Device Driver Errors	821
Application Install, Start, and Load Failures.....	821
Corrupted shortcuts	821
Damaged PIF.....	822
Corrupted memory space	822
Applications will not install	823
Solving Other Problems	823
General protection faults.....	824
Illegal operation.....	824
Invalid working directory	824
Optional device will not function	825
Terminate Stay Resident (TSR) programs and viruses.....	826
Getting an A+	827

Chapter 3: Preparing for Disasters with Disaster Recovery 831

Working with Boot Disks	831
Boot disks for Windows 9x.....	831
Windows XP	835
Using F8 Options during Boot-up.....	836
Booting into Safe Mode.....	836
Tracking the boot with a bootlog.....	838
Using the Last Known Good Configuration.....	839
Other Windows 2000 and Windows XP boot options	840
Entering MS-DOS Mode.....	840

Emergency Repair	841
Emergency Repair Disk (ERD)	841
Emergency Repair Process.....	841
Recovery Console.....	843
Understanding Backup and Restore	844
Backing up your computer.....	845
Restoring files from a backup	849
Restore points in Windows XP.....	850
ASR in Windows XP	853
Recovery and Rescue Methods	856
Recovery CDs and Recovery Partitions	856
Rescue CDs.....	856
Rescue Partitions or Rescue Installations.....	860
Getting an A+	861

***Book VIII: Networking*865**

Chapter 1: Down to the Networking Basics867

Identifying the Types of Networks	867
Peer-to-peer networks	867
Server-based (client-server) networks.....	869
Additional networking terminology	871
Understanding Network Topologies	871
Bus.....	871
Star	873
Ring	874
Hybrid	874
Wireless	875
Connecting with Network Cabling	876
Twisted pair	876
Coaxial	878
Fiber optic	882
Troubleshooting Networking Cables	884
Examining Network Access Methods	884
CSMA/CD	884
CSMA/CA	885
Token passing	885
Piecing Together the Network Architectures	886
Ethernet.....	886
Token Ring.....	888
FDDI.....	889
Understanding Network Protocols	889
IEEE Standards	889
Voice over IP (VoIP).....	891
Working with Network Devices	891
Network interface card (NIC).....	891
Repeater	892

Bridge.....	893
Router	894
Gateway	895
Hub	896
Switch.....	896
Wireless access point (WAP)	897
Understanding Communication Methods	898
Ways to Network a Computer	898
Network card.....	899
Serial and parallel ports	899
Infrared port.....	900
Getting an A+	900

Chapter 2: Understanding Common Wireless Communications907

Understanding Infrared Devices.....	907
Working with Wireless Networks	908
Components of a wireless network.....	908
Connecting to a network	910
Wireless standards	910
Securing wireless networks	914
Other Wireless	919
Bluetooth.....	919
WAN cellular.....	920
Increasing Wireless Network Performance.....	920
Troubleshooting Issues	921
Getting an A+	922

Chapter 3: Networking the Operating System927

Understanding Networking Components.....	927
Network adapter driver	928
Network client.....	932
Protocols	935
Services	938
The TCP/IP Protocol	939
IP address	940
Subnet mask.....	942
Default gateway	943
Configuring TCP/IP in Windows 2000/XP/2003.....	944
Configuring TCP/IP en masse using DHCP	945
Automatic Private IP Addressing	946
Understanding Name Resolution	947
NetBIOS names	947
Fully qualified domain names (FQDN).....	950
Troubleshooting with TCP/IP Utilities.....	954
IPCONFIG	955
PING.....	955
TRACERT	956
NBTSTAT	957

NETSTAT	958
PATHPING	958
NSLOOKUP	958
Sharing File System Resources	959
Share-level access control.....	959
User-level access control	960
Enabling File and Printer Sharing in Windows 2000/XP/2003	960
Creating shared folders	961
Hidden shares	966
Multiple shares	966
Connecting to shares	966
Sharing Printer Resources	968
Sharing a printer in Windows 2000/XP/2003.....	968
Installing a network printer in Windows 2000/XP/2003.....	968
Installing a network printer by using Point and Print	969
Understanding Windows Services	970
Restarting a Windows service	970
Server service.....	970
Workstation service	971
DHCP Server service.....	971
Print Spooler service	971
Messenger service.....	971
Browser service.....	972
Getting an A+	973

Chapter 4: Configuring Internet Access 981

Understanding the Internet	981
Using an ISP.....	983
Working with an ASP.....	985
Understanding Internet Protocols	985
TCP/IP	986
E-mail.....	987
HyperText Transport Protocol	989
File Transfer Protocol	992
Command shells	993
Installing and Configuring Browsers.....	994
Configuring General settings.....	995
Configuring Security settings.....	995
Configuring Privacy settings	996
Configuring Content settings	997
Configuring Connections settings	997
Configuring Programs settings	998
Working with Advanced settings.....	998
Using the Internet.....	998
Access methods.....	999
Domain names and Web sites	1000
Ways to Access the Internet	1002
Dial-up.....	1002
Cable	1002

ADSL	1002
ISDN	1003
T1/T3	1003
Satellite	1003
Wireless	1004
Firewalls and Home Routers	1004
Getting an A+	1005

***Book IX: Securing Systems* 1011**

Chapter 1: Fundamentals of Security1013

Identifying Types of Attacks	1013
Social engineering attacks	1014
Network-based attacks	1015
Software-based attacks	1019
Understanding Physical Security	1020
Understanding Authentication and Authorization	1023
Authentication	1023
Authorization	1025
Methods of Securing Transmissions	1027
Don't Forget about Data Protection	1028
Destroying data	1029
Backing up data	1029
Implementing RAID solutions	1031
Getting an A+	1032

Chapter 2: Implementing Security1039

Securing Systems through BIOS	1039
Implementing Users and Groups	1040
Creating user accounts	1040
Creating groups	1043
Implementing Permissions and Rights	1045
Rights	1046
Permissions	1048
Implementing Auditing	1051
Enabling auditing	1051
Reviewing the Security Log	1053
Implementing Firewalls	1054
How a firewall works	1054
Creating a DMz	1056
Implementing Security Best Practices	1057
Hardening a system	1057
Patching systems	1057
Firewalls	1057
Password policies	1058
Auditing	1058
Use switches instead of hubs	1059

Use anti-virus software	1059
Securing wireless.....	1059
Getting an A+	1060
Chapter 3: Viruses and Malicious Software	1065
My Computer Has a Virus!	1065
Types of viruses.....	1065
Protection from viruses.....	1067
Using virus-protection software	1068
Updating virus definitions.....	1070
Understanding Malicious Software.....	1071
Trojan horses	1071
Worms	1072
Spyware and adware	1072
Identifying Hoaxes.....	1073
Preventing Problems by Staying Up to Date.....	1074
Windows Update.....	1074
SUS and WSUS.....	1074
Antivirus	1076
Other security software.....	1076
Getting an A+	1078
<i>Appendix A: About the CD.....</i>	<i>1083</i>
<i>Appendix B: CompTIA A+ Exam Reference Matrix.....</i>	<i>1087</i>
<i>Index.....</i>	<i>1099</i>

Introduction

The *A+ Certified Professional* certification is a well-recognized certification and will serve as a basic foundation for a number of other certifications that you may eventually pursue. The certification exam tests your knowledge of both hardware and software used in today's computer world and the certification is one of the most popular certifications that IT professionals attain to prove their hardware and software knowledge.

About This Book

The *A+ Certification All-In-One Desk Reference For Dummies* is designed to be a hands-on, practical guide to help you pass the A+ certification exam. This book is written in a way that helps you understand complex technical content and prepares you to apply that knowledge to real-world scenarios.

We, the authors of this book, understand the value of a book that covers the points needed to pass the exam, but we also understand the value of ensuring that the information helps you perform IT-related tasks when you're on the job. That is what this book offers you — key points to pass the exam combined with practical information to help you in the real world, which means that this book can be used in more than one way:

- ◆ **As an exam preparation tool:** This book's key focus is to help you pass the A+ exam, and as a result is packed with exam-specific information. You should understand everything that is in this book before taking the exam, but to help identify key points that you *must* know, we use the *For the Exam* icon — watch for those as you read through the book in preparation for the A+ exam.
- ◆ **As a reference:** We have extensive experience in the IT industry and know the importance of ensuring that you not only can pass the exam but also can perform common computer-related tasks on the job — this is where this book's value as a reference tool shines through.

Conventions Used in This Book

Each chapter in this book has different elements that help you prepare for your A+ exam. The following sections can be found in each chapter:

- ◆ **Labs:** Throughout the chapters, lab exercises offer the opportunity to get your hands dirty with a particular topic. The labs are designed to give you real-world experience performing specific tasks. The labs for the chapters are located on the CD-ROM in a file called `Labs.pdf` in the Author directory.
- ◆ **Icons:** There are a number of different icons that are used in each chapter that draw your attention to information that is needed for the exam or in the real world. For more details on the icons used in each chapter check out “Icons Used in This Book.”
- ◆ **Getting an A+:** Found at the end of each chapter, “Getting an A+” covers key points you should remember for the exam.
- ◆ **Prep Test:** After the “Getting an A+” section, you are presented with example questions that help prepare you for the A+ exam. These questions are designed to be at the same level of difficulty as the A+ exam itself.

Foolish Assumptions

We make a few assumptions about you as a reader and have written this book with these assumptions in mind:

- ◆ **You are interested in obtaining the A+ Certification:** As mentioned previously, the key focus of this book is passing the exam.
- ◆ **You have a computer to work on:** In order to perform the labs that are located on the CD-ROM, you need a computer that you can take apart and install a copy of Windows on.
- ◆ **You have some experience using the computer:** As a prerequisite to the A+ certification, you should already feel comfortable with the computer. For example, you should know how to use the mouse and should know the difference between a left-click and a right-click. You should also know your way around Windows a little. For example you should know that the Start button is where different programs can be started. These are examples of tasks we consider “*using the computer.*”
- ◆ **You will study hard and do as much hands-on work as possible:** There is a lot of content on the A+ exam — you will most likely need to read over the information a few times to ensure that you understand the content. You should also experiment as much as possible after you read a particular topic. For example, after you read about how to add a hard drive to the computer, you should get a second hard drive and add it to your computer.

How This Book Is Organized

As with all of the *All-In-One For Dummies* books, the chapters are organized into minibooks. The chapters in each minibook are related by a specific theme or topic. For example, Book IX is called “Securing Systems” and contains all security-related content needed to pass the A+ exam.

The following sections outline what you can find in each minibook:

Book I: Setting the A+ Groundwork

In this book, you discover what A+ certification is all about and what you will be tested on when taking the exam. You also find out some basic safety guidelines and general “soft skills,” such as how to communicate with customers.

Book II: Inside the Box

In Book II, you find out about the inside of a computer and what makes it tick. You venture into the innards of your computer, such as motherboards, memory, and processors, and learn how to install and upgrade these components.

Book III: Outside the Box

Book III discusses topics related to what goes on outside the computer. For example, you find out about the different ports on the back of the computer and how to install devices such as video cards, sound cards, and network cards that give you output such as sound and video.

Book IV: Maintenance and Troubleshooting

The maintenance and troubleshooting minibook discusses how you care for the computer in order to prevent problems or hardware failure. This minibook also discusses how to troubleshoot different hardware components and how to diagnose the problem.

Book V: Operating System Basics

Book V is a special minibook because it makes the transition from hardware topics to operating system topics. In this minibook, you find out the purpose of an operating system and which major files an OS needs in order to run. You also discover how to install and upgrade Windows.

Book VI: Managing the Operating System

In this minibook, we discuss and demonstrate in a hands-on way how to manage aspects of the Windows operating systems. You find out how to load drivers and how to install and support applications and start using a number of troubleshooting tools available within Windows.

Book VII: Recovering Systems

In this minibook, you discover which core files are needed to start up a Windows computer and how to fix computers that fail to boot up. You also find topics related to system recovery, such as how to back up the system and perform a restore operation.

Book VIII: Networking

The A+ exam tests your knowledge of basic networking concepts and ensures that you know how to network two computers. This minibook gives you the networking background you need to pass the A+ exam. You discover the basics of networking technologies and figure out how to troubleshoot systems on a TCP/IP network.

Book IX: Securing Systems

This final minibook covers topics related to securing your environment. In this minibook, you go through the fundamental terms related to network security and how to perform tasks such as creating user accounts, setting permissions, performing virus scans, and updating Windows.

Appendixes

Appendix A gives you an overview of what you can find on the CD-ROM that accompanies the book. Please have a look at this section as there are valuable resources on the CD-ROM — such as lab exercises!

Appendix B is an exam objective mapping table that lets you know where in the book each of the exam objectives can be found. This is very useful when you are preparing for the exam and want to make sure you know each of the points in the objectives.

Icons Used in This Book

We use a number of icons in this book to draw your attention to pieces of useful information:



This icon gives you a heads-up on information you should absolutely know for the certification exam.



Information that would be helpful to you in the real world is indicated with a Tip icon. Expect to find shortcuts and timesavers here.



This icon is used to flag information that may be useful to remember on the job.



Information that could cause problems to you or to the computer are indicated with a Warning icon. If you see a Warning icon, make sure you read it. The computer you save may be your own.



Detailed information that is not needed for the exam or that is a step above the knowledge you absolutely need to know for the exam is indicated with a Technical Stuff icon.



This icon lets you know when you can find accompanying information or supporting documents on the CD-ROM. For example, there are a number of lab exercises written for the book — when a topic has a related lab, it is indicated with this icon.

Where to Go from Here


The A+ Certification is the most popular certification for individuals new to the IT industry and the certification world. After you have passed the A+ exam, you may want to continue your certification path by studying for the following certifications:

- ◆ **CompTIA's Network+ certification:** Network+ is a vendor-neutral certification that ensures the candidate understands networking technologies.
- ◆ **CompTIA's Security+ certification:** After passing the Network+ certification, you may want to continue on the network topic with a certification that proves your network security knowledge.

Book I

Setting the A+ Groundwork

The 5th Wave By Rich Tennant



"Wait! Wait! Wait! You've got a lung and two eyeballs in there! I thought you said you were A+ Certified!"

Contents at a Glance

Chapter 1: The New A+ Exams	9
Chapter 2: A+ Soft Skills	19
Chapter 3: Protecting Yourself and Your PC	35
Chapter 4: An Overview of System Components	49

Chapter 1: The New A+ Exams

In This Chapter

- ✓ **Understanding A+ Certification and its benefits**
- ✓ **Looking at the exams and their objectives**
- ✓ **Preparing for the exams**
- ✓ **Arranging to take the exams**
- ✓ **Test day**

So you are interested in taking the CompTIA A+ Certification exams? This chapter introduces you to the exams and gives you a good idea of what you can expect when you go to take them. Knowing what to expect in regard to the exam procedures and format will remove that uncertainty from you, which can weigh on your mind. Read through the procedures here; then you will be able to focus on the exam facts, which will help you breeze through the exam.

Hopefully, this chapter will remove some of that normal fear of the unknown you may experience by giving you information about the actual test-taking process. It also helps you develop good test-taking skills.

CompTIA A+ Certification and Why You Need It

What is the benefit of the CompTIA A+ Certification? It is proof to whomever you meet that you know and have validated the hardware and software knowledge that is necessary to support the troubleshooting and repair of computers. The CompTIA A+ Certification can be presented to employers and clients alike as proof of competency and skill in this area. This certification never expires.

CompTIA is a 22-year-old company that is focused on providing research, networking and partnering opportunities to its 19,000 members in 89 countries. In 1993, in response to the need for vendor-neutral, entry-level PC certification, the company created the A+ Certification. Prior to CompTIA creating the A+ Certification, there were many places a person could get hardware and software certifications; but this was often very expensive, difficult to get training, and not designed for accessibility for most people.

Microsoft, Novell, IBM, and other software companies offered software certifications, but these were specifically focused on teaching high-level support skills for these products, difficult for average users or support people to attain, and lacked relevance for most day-to-day work. IBM, HP, Compaq, SUN, and other hardware companies offered hardware repair and maintenance certifications, but again, these were specifically focused on their hardware and more on the peculiarities of their own platforms, and did not always cover the basics of configuration and maintenance. CompTIA stepped in to fill the gap that a majority of users fell into, which is a hardware and software neutral certification that covered all of the basics that are required by a support person. This certification can then be followed by vendor-based certifications, if desired; but the A+ Certification by itself proves a firm grasp of the basics.

An A+ Certification gives *employers* confidence that existing employees or new recruits have a level of knowledge that will allow them to do their jobs efficiently. It also gives employers a yardstick against which recruits and employees can be measured. A+ Certification also allows *clients* to rest assured that the person they hire to fix their computers has the knowledge to do so without blowing up equipment or deleting valuable data. This provides clients with peace of mind and increases repeat business. In the end, with the CompTIA A+ Certification on your side, you have more opportunities open to you in your career path.

Checking Out the Exams and Their Objectives

You have to take two exams to get your CompTIA A+ Certification. You have one required exam and one elective exam. The required exam is the CompTIA A+ Essentials exam, which focuses on the terminology and concepts. The elective exams are technician exams, each of which focuses on knowledge required to work in different IT environments. The technician exams have a very similar core, with differences based on the exams' particular focus.

The revised exams for CompTIA A+ were released in 2006. You will have 90 minutes to complete each exam and Table 1-1 contains the number of questions and passing score for each exam. CompTIA is releasing the exams as *linear format* exams or standard timed exams which are taken on a computer. After they have gathered grading statistics, they may re-release the exams as adaptive exams, which is what CompTIA has done in the past; but there are no current plans by CompTIA to do so. If the exam you take is adaptive, then these limits will change, as *adaptive* exams ask a minimal number of questions (usually about 15), and then ask additional questions based on any incorrect answers. The exam adapts to your wrong answers by choosing additional questions for you from the area where you are weaker.

<i>Exam</i>	<i>Number of Questions</i>	<i>Minimum Passing Score</i>
A+ Essentials (220-601)	100	675
IT Technician (220-602)	80	700
Remote Support Technician (220-603)	90	700
Depot Technician (220-604)	90	700

If you are writing an adaptive exam, you will pass or fail based on the number of wrong answers you get. For instance, if on a standard timed exam you are asked 50 questions and you need a score of 75% to pass, you need to get 38 correct answers. For the same exam in an adaptive form, you would fail if you get 13 questions wrong. As soon as you get to that level, your test would be over, even if you were asked only 13 questions, since after that point, it is then statistically impossible for you to pass. After an initial set of questions, if you have few enough wrong answers, you will pass.

The CompTIA A+ Essentials Exam

The CompTIA A+ Essentials Exam (220-601) covers the basics of computer maintenance and support across the CompTIA A+ *domains*, or exam areas. Table 1-2 provides a breakdown of the exam areas that are covered on the CompTIA A+ Essentials Exam. This exam puts heavy emphasis on personal computer components and operating systems.

<i>Domain</i>	<i>Percentage of Examination</i>
1.0 Personal Computer Components	21%
2.0 Laptop and Portable Devices	11%
3.0 Operating Systems	21%
4.0 Printers and Scanners	9%
5.0 Networks	12%
6.0 Security	11%
7.0 Safety and Environmental Issues	10%
8.0 Communication and Professionalism	5%

In addition to the CompTIA A+ Essentials Exam, you will have to take one elective exam. You can choose from these three other exams to complete your CompTIA A+ Certification:

12 *Checking Out the Exams and Their Objectives*

- ◆ The CompTIA A+ IT Technician Exam (220-602)
- ◆ The CompTIA A+ Remote Support Technician Exam (220-603)
- ◆ The CompTIA A+ Depot Technician Exam (220-604)



CompTIA has announced that they will be creating another elective exam, 220-605, but have not announced what it will be called or what it will cover. Likely, this additional exam will be developed based on feedback on the current exams to respond to an unanticipated need. You can expect that it will have a slightly different focus, but share the same core as the existing exams. The contents of this book will prepare you for this new exam.

CompTIA A+ IT Technician Exam

The CompTIA A+ IT Technician Exam (220-602) is geared toward individuals who work or will be working in field service or enterprise environments, such as service technicians, field PC technicians, IT technicians, IT administrators, ICS specialists, and computer technicians. The breakdown of the exam components is covered in Table 1-3. Based on the domain breakdown, this exam has the widest breadth of topics for a well-rounded IT technician and includes questions on laptops and safety.

Table 1-3 **CompTIA A+ IT Technician Exam (220-602)**

<i>Domain</i>	<i>Percentage of Examination</i>
1.0 Personal Computer Components	18%
2.0 Laptop and Portable Devices	9%
3.0 Operating Systems	20%
4.0 Printers and Scanners	14%
5.0 Networks	11%
6.0 Security	8%
7.0 Safety and Environmental Issues	5%
8.0 Communication and Professionalism	15%

The CompTIA A+ Remote Support Technician Exam

The CompTIA A+ Remote Support Technician Exam (220-603) is geared toward individuals who will remotely interact with customers for support purposes via the telephone or e-mail. People who take this test hold positions such as remote support technician, service desk technician, support center technician, Help Desk technician, and Call Center technician. The breakdown of the exam components is covered in Table 1-4. Based on the domain breakdown, this exam reduces the domains, but increases focus on operating systems, security, and communication skills.

Table 1-4 **CompTIA A+ Remote Support Technician Exam (220-603)**

<i>Domain</i>	<i>Percentage of Examination</i>
1.0 Personal Computer Components	15%
2.0 Operating Systems	29%
3.0 Printers and Scanners	10%
4.0 Networks	11%
5.0 Security	15%
6.0 Communication and Professionalism	20%

The CompTIA A+ Depot Technician Exam

The CompTIA A+ Depot Technician Exam (220-604) is geared toward individuals who work or will work on a repair bench with limited customer interaction or in scenarios that emphasize computer repair and troubleshooting.

Depot repair technicians and bench technicians benefit from this exam.

Table 1-5 shows the breakdown of the exam domains. Based on the domain breakdown, this exam increases focus on the components of computers, and the hardware areas of A+ Certification.

Table 1-5 **CompTIA A+ Depot Technician Exam (220-604)**

<i>Domain</i>	<i>Percentage of Examination</i>
1.0 Personal Computer Components	45%
2.0 Laptop and Portable Devices	20%
3.0 Printers and Scanners	20%
4.0 Security	5%
5.0 Safety and Environmental Issues	10%

Using This Book to Prepare for the Exams

Exams are stressful events for most people, but if you are well prepared your stress level should be much lower. If you read and understand the material in this book, you should have no problem with any one of the exams. The review questions, sample exams, and exam test engine on the companion CD are all designed to prepare you for what lies ahead. There is substantial overlap between the topics covered on the CompTIA A+ Essentials exam, and the three technician's exams, so this book takes a holistic approach to studying for the exam. You should review all the material that is found here, and then you will be prepared to go and write any of your exams you choose to. It will be best to write both of your exams in a short timeframe, to avoid forgetting the information which you have already learned.

You should examine the objectives for each chapter before diving into the content so that you can use them as a guide to which sections you may need to focus on. After thoroughly reading the content of the chapter, attempt the Prep Test section at the end of the chapter. If you do poorly on the Prep Test, then go back to the objectives again to see where you need more effort. When you re-read the chapter sections, try to examine the content from another viewpoint since this may help you associate the information with the questions and the objectives. A differing viewpoint might be that of a computer user, a help desk employee, or a desktop support technician. After you can complete the Prep Test in each chapter with an 85% or better score, then you should attempt the exams on the companion CD. If you are unable to achieve a mark of 85% or better, then you should continue to review the areas in which you are weak.

Making Arrangements to Take the Exams

The A+ Certification exams can be scheduled at Pearson VUE or Thomson Prometric testing centers. For more information about scheduling your exam, check the CompTIA A+ Certification page on CompTIA's Web site at <http://certification.comptia.org/a/default.aspx>.

The cost of taking the A+ exams is US\$120 per exam for CompTIA members and US\$153 per exam for nonmembers. To become a CompTIA member, you must hold a valid CompTIA certification, and pay a membership fee. If you are reading this book, you probably do not have your A+ Certification, but you might hold a Network+, Linux+, or some other CompTIA certification. If so, then you will be able to become a CompTIA member. Details on this process can be found on CompTIA's IT Pro Web site at <http://itpro.comptia.org/default.aspx>.

The Day the Earth Stood Still: Exam Day

Knowing what to expect on the day of the exam can take some of the pressure off of you. The following sections look at the testing process.

Arriving at the exam location

Get to the exam location early on the day of the exam. You should arrive at the testing center 15 to 30 minutes before the exam starts. This keeps you from being rushed and gives you some temporal elbow room in case there are any delays. It is also not so long that you will have time to sit and stew about the exam. Get there, get into a relaxed frame of mind, and get into the exam.

When you get to the test site, before you sign in, take a few minutes to get accustomed to the testing center. Get a small drink of water. Use the restroom if you need to. The test will be 90 minutes, so you should be able to last that long before another break. You may want to check with the policy for bringing a drink of water in with you; some centers will allow it, while others will not.

Now relax. Getting to the exam site early gives you this privilege. You didn't show up early just to stew and make yourself more nervous.

If you feel prepared and are ready to go, you may want to see whether you can start the test early. As long as the testing seat is free, this is usually not a problem. When signing in, you will usually need two pieces of ID, with at least one of them being a photo ID. After signing in, you will be taken to the testing room.

Taking the exam

In the testing room, and depending on the size of the testing center, there will usually be from one to eight computers set up. Each computer represents a testing seat for this computer-based exam. The exam consists of multiple-choice questions. Take it slow, or at least pace yourself. Trying to complete the questions too quickly will no doubt lead you to errors. When you are about to start the exam, you will see on the screen how many questions there will be, and how long you have to complete it. Be sure to read the exam instructions on the screen at the start of the exam, as they do change from time to time. Based on the number of questions and your exam time, figure out how long you can spend on each question. On average, you have slightly over a minute per question. Take your time, but be aware of your time for the exam overall. When you have completed 25% of the exam, you should have used only 25% of your allotted time.

Read the entire question and try to decide what the answer should be before looking at the answer choices. In most cases, you will find a few key words that are designed to remove any ambiguity in the question, as well as a few distracters and useless information designed to throw you off. If you do not notice these key words, the question will seem vague. If this is the case, re-read the question and look for the key words. Exam questions are written by many authors, so the style of writing for each question may differ.

Don't overcomplicate the questions by reading too much into them. Besides the key words and the distracters, the question should be straightforward. In some cases, the question may ask for the best choice, and there may be more than one answer that seems correct. Choose the one that is *best* — quickest, most likely to succeed, least likely to cause other problems, whatever the question calls for. The best choice is always the right choice.

After identifying the key words and distracters, follow these additional steps:

1. Eliminate choices that are obviously wrong.

Most questions will ask you to choose one of four answers, while some will ask you to choose all that apply and have as many as eight choices. You should be able to immediately eliminate at least one choice — perhaps two. Now the odds of choosing the right answer have gotten substantially better. Re-read the question and your remaining choices carefully and you should be able to locate the correct answer.

2. If you don't have a clue which of the remaining choices is correct, mark an answer.

On a standard timed exam, you will be able to come back to review your answers. Not answering a question is automatically wrong, so if you have an answer, it might be right. You may also find information on other questions in the exam that triggers the correct answers for questions you were not sure of.

3. Make your choice and leave it.

Unless you have information that proves your choice is wrong, your first instinct is usually correct.

The adaptive exam delivers a series of questions to you. If you answer a question wrong, you get additional questions in that category. You can't review or change your answers on the adaptive exams, as you can with standard timed exams. Since skipped questions are automatically wrong, with no ability to change them, you must provide an answer to the question before moving to the next question. You are allowed to attempt to answer a finite number of questions (known only by CompTIA) before the test "decides" that you really don't know what you need to know to pass the test. If you exceed the allowed number of wrong answers (again, known only by CompTIA), the exam ends; you fail. (But that won't happen to you because you bought this wonderful book!)

If you get all of the initial category questions correct, you can pass the exam in relatively few questions, but if you get an initial category question wrong in a category in which you are weak, you will find the adaptive format very difficult. You won't know exactly how many questions you have to answer to complete the exam, so the end of an adaptive exam will always come as a surprise — hopefully a good one.

If you are taking the non-adaptive exam, you are allowed to mark questions and come back to them later. However, it is a good idea to select an answer for every question, even if you are unsure about it, because you might run out of time before you can review previous questions.



Remember that your first choice is usually correct — don't second-guess your first choice. Change your answer only if you're absolutely positive it should be changed.

Regardless of which type of exam CompTIA has available for you when you take your exam (adaptive or standard timed), you are given a Pass/Fail mark right on the spot after completing the exam. In addition, you get a report listing how well you did in each domain. If you don't pass (or even if you do), you can use this report to review the material on which you are still weak.

How does CompTIA set the pass level?

CompTIA uses a scale score to determine the total number of points that each question on the exam will be calculated out of. Your final score will be between 100 and 900. In any case, the passing score (not a percentage due to the scale) varies from one exam to the others. The scale score system allows the number of points assigned to questions to vary between each copy of the exam, which makes it harder for test candidates to compare scores across exams.

If you do not pass, then CompTIA has a retake policy. If you do not pass on the first attempt, you can take the exam again at any time; but you have to wait for at least 30 days before your third or subsequent attempt.

Chapter 2: A+ Soft Skills

Exam Objectives

- ✓ **Knowing your troubleshooting techniques**
- ✓ **Communicating effectively and professionally**

As a CompTIA A+ Certified Professional, you will often need to troubleshoot systems as well as deal with customers or clients. In order for you to complete these tasks, you not only need to have good troubleshooting skills, but you also need to develop good customer-relation skills.

This chapter examines the basics of troubleshooting; these basics are not based on specific hardware components, but rather basic methodology which will be used to solve problems. This methodology will not change, even though the underlying technology may change. Troubleshooting of specific components is covered in Book IV, Chapter 2.

In the process of solving problems for customers, be they internal customers or external customers, you will likely need to talk to them. This chapter will take a look at communication skills that you will want to use to help you maintain a high level of professionalism.

Using Troubleshooting Procedures and Good Practices

Many people say that troubleshooting skills cannot be learned, but this is not true. If you are methodical or just very curious, you will have an easier time troubleshooting, but anyone can learn the practical order that should be followed when troubleshooting problems. If you follow the steps in this chapter, you should have an easy time when it comes to troubleshooting.

There are two main sections to troubleshooting: hypothesis generation and hypothesis evaluation. In other words, based on the symptoms that you see, you need to create a list of things that might be wrong and then test to see whether any of them are actually the problem. Figure 2-1 shows a rough outline of general troubleshooting procedures.

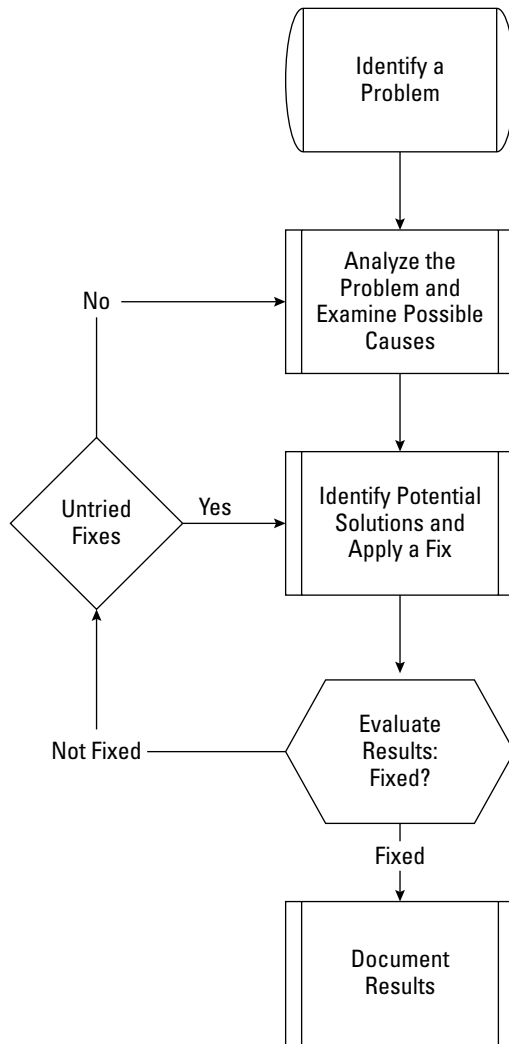


Figure 2-1:
Basic
troubleshoot
ing steps.



On the exam, you can expect that some of CompTIA's terminology may vary from what I use here, but the general guidelines and approximate order will remain the same.

In order to be able to troubleshoot any system, you need to know how that system is supposed to function and what options can be changed or might fail. This information allows you to quickly assemble a list of hypotheses. The better you know the system, the easier this phase will be. You then need to test these hypotheses. Testing may involve using tools appropriate to the

system to test the condition of certain components, or it may involve changing a series of configuration settings and re-checking to see if the problem has been resolved at certain stages.

Identifying the problem

The first thing you need to do is identify the problem. The first step to identifying the problem is questioning the system's user. This shouldn't be a police interrogation, nor should you accuse the user of causing the error. You might discover, in the end, that the user actually *did* cause the error, but you should never assume that up front or you'll quickly become an adversary instead of an advocate.

The information you gather from the user helps you set the scope of the problem. Many users start with the simple, "It's broke!" It is up to you to figure out how it is "broke," and you do that by asking simple questions: "What issues are you having with it? What is it not doing? Does it always happen or only at certain times?"

Even if the user gives you a better starting point, you will likely still need to ask some questions to refine the scope of the problem. Here are some things to keep in mind as you try to identify the problem:

- ◆ **Identify changes to the computer:** Some errors are caused by the user or his or her surroundings. The user should always be asked up front, in much the same way a doctor will ask if his lifestyle has changed or if he is on any medication. Ask the user if he has made any changes to the computer by installing software, changing hardware, or changing the computer's location. This will often provide you with an early lead to the possible problem.
- ◆ **Break big problems into smaller pieces:** If the problem is large, try to break it up into smaller sections. Going back to the doctor analogy, the "I feel sick" statement causes the doctor to ask questions and build a list of symptoms. That list of symptoms (sore throat, headache, runny nose, gaping hole in abdomen) may be easier to deal with if it is broken down in to groups of symptoms, rather than being treated as a whole. The doctor may then treat each group of symptoms individually. The same applies to the computer problems. That "broke" computer may be malfunctioning due to a variety of smaller problems, such as insufficient RAM, a fragmented hard drive, malware, Trojans, and poor networking protocol configuration. If the problems are broken down instead of being treated as a whole, you can approach each item separately.
- ◆ **Prioritize problems and decide in which order they should be tackled:** With your list of problems, you can decide which are the most critical and should be approached first. For instance, if the hard drive is

corrupted *and* you need to reinstall or upgrade the operating system, you should fix the issue with the hard drive before performing the software upgrade.

In most cases, you will prioritize the issues based on the impact to the user and the impact of other fixes implemented at a later time. One rule I live by is “fix the business problem first.” For example, if you are working for a stock trader, she will need her stock trading program working immediately, but she might be willing to wait for a fix to the word processing program.

Analyzing the problem and potential causes

For each problem that you work on, you need to decide the scope of the problem. This can be done by testing. The initial tests that you conduct will determine the scope of the problem. For example, at the doctor, if your stomach is upset, you will be asked stomach- and food-related questions, and then some basic tests will be conducted to determine whether the problem is an ulcer, an allergy, or just some bad sushi.

The same is true with computer issues. You ask questions and then perform rudimentary tests to determine the scope of the issue. For example, after being told that a printer does not print, you might verify that the cables are in place, that it powers up, that it has toner or ink, and that it accepts commands from the computer. These rudimentary checks allow you to see how large the problem is, and from there you can perform additional analyses. For example, if the printer is passing paper but the pages come out blank, then your list of potential causes can be reduced because it is not related to power. Rather, the problem is either with the transfer of the printing material — the ink or toner — or with the interpretation of the commands from the computer. With this narrowed list of potential causes, you can move on with addressing the problem.

The basic troubleshooting process

With your list of potential causes in hand, you can move on to figuring out potential solutions. In some instances you may approach the problem starting at a general level and then working down to specific components or you may look at entire systems, like the power system, in a sequential manner. In most cases, your potential causes can be placed on an imaginary line which forms a resolution path. The resolutions path forms a sequential chain of possible causes for the issue, starting at one point and moving to another. For instance, if the computer does not power up, you can start testing at the electrical socket, then move to the power supply, the power supply connectors, the power switch, the motherboard, and then to the devices. This process would move through the sequential chain along the possible path that power would flow. Following a possible resolution path from one end to the other is sometimes referred to as the *layered* or *linear approach*.



The term *layered* is often used when discussing networking issues, since the layers in the networking model are easily definable and are often referenced. For more information about the layered network model, refer to Book VIII which discusses networking.

Most systems that make up computer systems and networks can be broken into layered models, and troubleshooting can move through these layers. Another example of layered troubleshooting would be a gasoline engine, which has two major systems, the electrical system and a fuel supply system. The electrical system would have layered troubleshooting starting at a power source, such as a battery, through power supply cables, a distributor, and finally to a spark plug. The fuel supply system would start with fuel storage, supply lines, pump, filter, and end at a combustion chamber. In the event of a problem, both of these systems can be reviewed and tested from end to end, in a linear or layered manner, until the problem is located.

The following sections take a close look at some suggestions for implementing potential solutions.

Back up any settings or files that you need to change

You don't want to introduce more problems to a system than it had when you started. Before changing any BIOS or application settings, you should record the original settings or back them up if that facility exists. For instance, if an application that you are troubleshooting stores its configuration in a text or binary file, copy that file so you can quickly restore it to the initial state if your solution doesn't work.

Some computers allow you to make a backup copy of the CMOS configuration or back up the existing BIOS prior to overwriting it with an upgrade. Some things cannot easily be reversed, like the application of a Service Pack, so care should be taken when doing those tasks, such as setting a restore point or making a full system backup.

Test one configuration change at a time

This step is required to know what the fix actually is. If you change five things and the problem component now works, you don't know which change made the difference. In some cases, there may be a monetary cost. For instance, if the computer starts to blue screen on boot and you suspect a hardware fault, you could change the disk controller, hard drive, RAM, CPU, and power supply. After that, if the problem no longer exists, what was the solution? Are you going to replace all five components the next time you have the issue? By substituting one item at a time and testing at each step to re-create the issue, you will know which change you implemented corrected the problem.

Don't discount the obvious

When I get the message “Missing Operating System” or “NTLDR not found,” the first thing I do is look for a floppy disk that was left in the system. For power problems, is it plugged in? For network problems, is the cable plugged into the correct jack? These are simple first steps, but are often overlooked by the user. Just because it is obvious to you does not mean that it is obvious to the user. If the solution is not obvious, then go back to your layer approach, since it will not let you down.

Research to find potential fixes

If you run out of ideas for a specific problem, do a little research. The “Documentation resources” section, later in this chapter, gives you some ideas about where to begin.



There is nothing wrong with you, as a CompTIA A+ Certified Professional, not knowing the solution to a particular issue with a specific product, but it is important that you know where you can find the answers quickly.

Prioritize potential fixes for the problem

You have spent time researching the symptoms, and now you have a long list of items to check and possible solutions to apply; you must prioritize them. There are many ways that you may choose to do this. For instance, you can put the most likely fix at the top of your list, or you can put the quickest and least disruptive fix at the beginning of the list. In most cases, you will likely do the latter.

For instance, a good fix for many software-related problems is to re-image the computer or re-install the application. Although this often fixes the problem, it can be time consuming for you and disruptive for the user. Re-installation often involves checking for local files that need to be backed up, backing them up, verifying application settings, restoring the computer operating system, restoring backed-up user files, and re-applying settings — this usually leaves a lot of settings that the user needs to reset when she runs the application(s) for the first time after re-installation. Because this is time consuming and disruptive, it should be left toward the end of the list of options. Rather, the list should begin with items like verifying application settings, testing new-user profiles, and other quicker and less disruptive fixes.

Test related components

Sometimes a symptom in one area is caused by a problem in another. The classic example with computers is being asked to look at a computer that is reported as slow, and has an extremely active hard drive. You can test the

hard drive and related cables, use diagnostic utilities to test its health, swap it with another, and not find a solution, only to take a step back and evaluate other system components such as memory. In many cases, a lack of RAM can cause excessive paging which causes high disk activity. So the solution to the problem is to add RAM, and to not do anything with the hard drive. If you did not look at the related component (RAM), you would not be able to solve the problem with the active hard drive. In some cases you may need to do some research to know what components of the system are related. Don't just go for the obvious; test related systems as well. That may be where the actual issue is.

Evaluate results

At each step of the analysis and application of potential fixes, you should evaluate the results. By evaluating often, you will know what is working, what is not, and what additional problem you may have caused. At any time, the results may cause you to change your initial hypothesis and choose to pursue a different path. Changing your hypothesis should be done only if there is sufficient proof that your initial statements were in error, or you may find yourself running in circles.

Document findings, activities, and outcomes

The philosopher George Santayana said, "Those who cannot remember the past are condemned to repeat it." In order to remember it for the long term, you must write it down or otherwise record it. If you do not document your findings, then you may find yourself repeating the same troubleshooting steps in the future to solve a similar problem, when the process could have been abbreviated.

It is a good idea to have a personal or corporate knowledge base to record troubleshooting activities. Some companies have implemented a formal help desk or trouble ticket system with the main goal of documenting their problems and solutions. Having a technician close his trouble tickets with "Problem resolved" may be fast for the technician, but it is not the purpose of the application. One of the application's goals is to build a database of problems and resolutions. The more detail that is included in the resolution documentation, then more helpful it will be in the future. Once the trouble ticket database has enough information in it, it becomes more than a place where the information is stored; it also becomes a place where solutions are found. The larger the database, the more likely it will be able to provide your solution. When a problem is reported, you will be able to search the database to see if the problem has previously happened for that person, or another person, and you may be provided with a solution. Testing a few known solutions is faster than starting the troubleshooting process at the beginning.

Documentation resources

CompTIA A+ Certified Professionals aren't expected to know every answer, but they should have a good list of resources at their disposal to find an answer. Here are the most common resources:

- ◆ **Vendor manuals:** Gone are the days of an IT department filled with bookcase upon bookcase of manuals for every application that the company owns. In many cases, manuals are now shipped on CD in PDF format or are available for download from vendors' Web sites. Many of these manuals have specific sections detailing product installation, options to watch during the installation, and diagnostics or troubleshooting. Looking up a computer's symptoms in this guide can lead to concise, quick solutions.
- ◆ **Service documentation:** In addition to general user manuals, some companies offer additional service documentation that can be obtained by owners of the product. In some cases, this will be available to you only if you are certified on the product or you are employed by the manufacturer. It is worth checking to see if this technical documentation is available to you, as it often has detailed configuration and troubleshooting information.
- ◆ **External knowledge bases:** In the same way that your organization can create its own internal knowledge base of problems and solutions, many vendors have compiled their own, and they make this data available and searchable on their Web sites. If the manuals seem to be lacking, you should see if vendor has a searchable online knowledge base. In some cases, vendors ignore the manual when they maintain a good knowledge base. Again, if your problem is found in the database, then you likely have a quick solution that you can apply.
- ◆ **Training:** Knowledge is half the battle, and one of the best ways to get quick knowledge is to attend vendor training on the products that you are working on. In many cases, you will find out a product's "gotchas," those common but hard-to-solve problems that are unique the product. If you don't walk away with this, then you should at least leave the training with a complete picture of how the product is suppose to work, which can reduce your troubleshooting time.
- ◆ **Other Internet resources:** In addition to the vendor Web sites as an official support tool, additional discussion forums and Web sites may have answers to your specific problems, even though the Web site or forum is geared toward a general industry area. For instance, you may be having issues with a specific sound card, and if the vendor's Web site does not provide you with the answer, you may find the answer on a site that covers all computer audio cards and issues. Your favorite search engine is a good place to start looking for these sites.

With these documentation resources and the basic troubleshooting guidelines that have been discussed, you should be able to approach problems and troubleshoot them in a methodical, organized way.



Additional troubleshooting processes for specific hardware components are covered in other sections of this book, with much of it in Book IV, Chapter 2.

Professionalism and Communication

Whether you are an internal or external support professional, you have customers. For the external support professional, it is easy to identify who that is. For the internal support professional, the people for whom you provide support are still your customers, and you still need to provide them with service. In order to provide this service, you will be required to communicate with them, and this communication should remain professional in its tone and manner. The best way to maintain the right level of professionalism is to respect the customer; whether the customer is right or wrong, you should still give him respect. You can still show him respect while you tell him that he is wrong.



The professionalism and communication section of the A+ exam should be common sense for most readers. The exam questions will present a wide range of responses to given scenarios; the answer will usually be the one that illustrates good communication skills and respect for the client. Verbally abusing the client and hanging up will not be the correct answer!

Good communication skills

What we've got here is failure to communicate.

—Cool Hand Luke

A breakdown in communication can be blamed for many of the problems that exist in our world today. It is often a source of conflict that begins with a misunderstanding. In order to avoid this conflict with your customers, you should always make sure you use good communication skills. Here is a list of things to keep in mind when communicating with a customer:

- ◆ **Listen to the customer and don't interrupt.** One of the best abilities of a good communicator is the ability to listen. All too often, in today's society, the emphasis is on speed, and for the sake of speed, people tend not to listen. Part of "not listening" is completing other peoples' thoughts or sentences. This is a problem that can be resolved by resisting the urge to just jump in. Put yourself in the listener's position while you let the

person tell you the problem and how it is affecting her life. “Not listening” can extend to e-mail. I have written e-mails to vendor support addresses, detailing the problem and the steps already taken to solve the problem, just to receive a stock reply from the vendor support person saying “Try this,” which was already tried, or does not address the problem. At times like that I usually reply to the e-mail, stating the fact that he or she failed to read message; but rather suggested a solution based on the subject line or the first line of the message.

- ◆ **Clarify statements and repeat what you have heard.** One technique that helps to clarify any misunderstandings is to repeat to the user what he has told you. This reinforces for the user that you are listening and are not daydreaming. It also verifies that you understand what the user said. When you repeat it back, you may substitute terms or give additional examples of what the client described to you.
- ◆ **Use clear, concise, and direct statements.** Translation: Don’t be long-winded. Many people like to hear themselves talk; resist any temptations to verbosity you might have. Don’t use the ten-dollar word if a twenty-five-cent word will do. Keep your statements direct with clearly focused points.
- ◆ **Avoid tech talk.** As a CompTIA A+ Certified Professional, you are constantly exposed to technical terms and abbreviations that you slowly incorporate into your vocabulary. However, not everybody in the general public is familiar with these terms. *RAID*, *asynchronous*, *MTBF*, *SCSI*, and a whole slew of other terms can confuse clients who refer to the computer as being composed of a monitor and a hard drive. Whenever possible, try to gauge the user’s knowledge level and use an appropriate level of terminology.

Professional behavior

You should always try to conduct yourself with professionalism. Doing so generates respectability in the eyes of your clients and supervisors. The following list contains some aspects of professionalism that should be practiced:

- ◆ **Positive attitude:** It is often said that animals can smell fear on a person. The same is true of users: They can smell a negative attitude right away. You should approach each client and each job with a positive attitude and a feeling of confidence. If you do, you will get more satisfaction from your job. Regardless of your position, there will likely be some aspects of your job that you don’t enjoy, but don’t let these things bring your attitude down.

In most cases, you deal with clients who have problems. These problems may put the client in a foul mood. If you seem happy to see a client, then he will likely lighten his mood, and you will both have a more positive experience.



This could go in reverse if you let it: The client's negative attitude could dampen *your* mood. Don't let this happen. Stay positive.

- ◆ **Proper dress:** Professionalism likely includes appropriate attire as well as attitude. When choosing what to wear, you should plan around the jobs that you are doing. Client meetings might dictate formal business attire, while if you are expecting to pull cable through crawlspaces, then something more casual would fit the task.



Regardless of the job, your clothes should be clean and well-mended. To a client, a shoddy appearance indicates shoddy work.

- ◆ **Privacy, confidentiality, and discretion:** It's surprising what information clients are willing to let you see on their computers. In many cases, you will have access to their most private documents. In corporate environments, this may include confidential employee information, and when doing retail bench work, this may include private financial information. In both cases, a client's computer often holds information that you should not be seeing and that, in normal circumstances, you would not be allowed to see. In most cases, you will know what you should and should not be looking at. If there is a reason to open a file to verify accessibility, then you will need to do that, but you should ignore the information you find when you're done.



As an outside consultant, I am amazed by the number of clients who freely offer up administrative credentials for their network, granting access to all areas of the network and to all kinds of information. Take care to ensure that these credentials are not misused by you or put in a place where they could be used by another person without your knowledge.

- ◆ **Respect:** People with different cultural backgrounds will have different means of showing respect, and you should use techniques appropriate to your local region or culture. In general the techniques listed here are fairly universal. Respect should be thought of as a two-way street. In most cases, if you show appropriate respect to your clients, they will return that level of respect.

If the client doesn't show you any respect, don't lower yourself to his or her level — stay calm and continue to be respectful. Ask the client to calm down, let him or her know that he or she is not acting appropriately, and that you will not be able to help if he or she does not give you the information that you need. This is a tough situation to call, as when people are told that they are not acting appropriately, they will either "hear" you and calm down, or they will become defensive and more disrespectful. As you tell them that they should calm down, make sure to reinforce that you want to help them, and the logic of the situation may finally sink in.

Some steps or actions you can take to show respect are:

- **Don't argue or be defensive.** The client may be in a foul mood. You would likely be as well if technology got in the way of your job or you thought that you had just lost several years of personal data or photos. A client may have many reasons to be upset when you start dealing with him.

To some clients, you represent the technology they're having issues with. Their frustration will often get vented all over you. Don't take this personally or get upset by users' actions. Do what you can to relieve some of their stress, and remind them that you're not part of the problem, you're part of the solution.

If you argue with the client, the entire issue will become ugly. Avoid arguing at all costs.

- **Don't belittle or minimize problems.** In most cases, the problems that users face are serious to them. Even though you may not share their view of the importance of the problem, do not dismiss the issue. By claiming that the problems are trivial, you suggest that the clients' feelings are trivial. If the clients were not upset before this point, they will be if they think you are trivializing them.
- **Don't be judgmental.** You might personally possess many biases. These should be checked at the door when you come to work. By leaving that baggage out when you arrive at work, you will likely reserve your judgments on issues until you have seen all of the facts. This point goes very nicely with the previous one — if you have passed early judgment on a user, you may be more inclined to trivialize her issues.
- **Avoid distractions and interruptions.** When dealing with a client, try to reduce distractions for yourself. If you are working a help desk, you should not be carrying on conversations with those around you when you are on a support call. If you are working a retail repair bench, you should deal with one client at a time and not try to multi-task among them or take phone calls. If you are visiting a client, then you should silence your mobile phone.

In all cases in which you are physically with a client, maintain eye contact. This will be taken as a sign that you're paying attention. If you don't maintain eye contact, the client will think that you are not paying attention. There is also a chance that you may become distracted by something you see. If you maintain eye contact, you will not get distracted by other events, and you will tend to have a better understanding of what is being said to you.



- **Property is as important as people.** Respecting the person extends to respecting his or her property and privacy. For many clients, you will be working in their office, very possibly at their desk. When you're waiting for a reboot or for software installation, you will be tempted to look around. Don't do it! You should always respect your client's privacy and possessions, from the computer to the "Marvin the Martian" action figure.
- **Don't bad-mouth past clients:** Respect for a customer extends beyond the time you spend fixing his computer. Don't be disrespectful of previous clients in front of current ones. If the client sees your lack respect for others or hears you say negative things about other clients, then she will think that you are going to say similar things about her to your other clients.

Getting an A+

In this chapter, I discuss general skills that will make you an effective troubleshooter, as well as how you present yourself to and treat the clients that you deal with. The following points are covered:

- ◆ Create a list of possible problems based on the information that is provided to you, and then follow that up with testing to locate the true cause.
- ◆ When troubleshooting problems, work through systems sequentially from general to specific or from beginning of a system to its end. That way, you won't miss items that may be involved in the solution.
- ◆ Back up user or system files and any configuration files or settings that may be lost or changed during troubleshooting.
- ◆ Document issues, solutions, and any other items that are important.
- ◆ Follow basic principles for respecting other people, which will answer most exam questions related to professionalism and communications.

Prep Test

- 1 Which of the following is not a step in the basic troubleshooting process?**
 - A Document findings
 - B Respect client property
 - C Evaluate results
 - D Identify the problem

- 2 Where should you go to get a complete list of problem symptoms?**
 - A Product documentation
 - B Internet search engines
 - C Question the user
 - D Analyze the problem

- 3 Which of the following should you find out when questioning the user?**
 - A Recent changes to the computer
 - B The last time the computer was shut down
 - C How the computer is connected to the building electrical system
 - D The location of resources to document your outcomes

- 4 When making configuration changes during the troubleshooting process, what practices should be followed? Choose all that apply.**
 - A Apply one configuration change at a time
 - B Discount the obvious
 - C Back up any settings or files that you need to change
 - D Test related components

- 5 When looking for information related to the problem, which of the following is not considered to be a good source? Choose all that apply.**
 - A Vendor manuals
 - B Service documentation
 - C Internet Web pages related to the product
 - D Competitive knowledge base

- 6 Which of the following is not a good communication skill?**
 - A Let the customer talk; don't interrupt
 - B Have a positive attitude
 - C Make concise statements
 - D Avoid jargon

Answers

- 1** B. Respecting client property is part of good communication skills, not the troubleshooting process. *See "The basic troubleshooting process."*
- 2** C. The best source for finding out what the symptoms are, and a description of what appears to wrong with the computer, is directly from the user. *Review "Identifying the problem."*
- 3** A. The user should be questioned about changes that have been recently made to the computer. In some cases a user may not know or may not want to tell you. *Check out "Identifying the problem."*
- 4** A, C. Before applying any changes to the system, relevant files and settings should be backed up so that the changes can be reversed or critical files can be restored if anything goes wrong. You should apply only one change at a time, so that you will know what change actually resolved the problem. While testing related components is part of the overall troubleshooting process, it does not necessarily apply to making configuration changes. You should not discount the obvious. *Peruse "The basic troubleshooting process."*
- 5** A, B, C. When looking for an answer related to a vendor's product, you might find some information on a competitor's knowledge base, but this is not typically going to be a good source of reference for you. *Take a look at "Documentation resources."*
- 6** B. Having a positive attitude is part of professional behavior and not a communication skill. *Peek at "Professionalism and Communication."*

Chapter 3: Protecting Yourself and Your PC

Exam Objectives

- ✓ Gathering the tools of the trade
- ✓ Understanding ESD and EMI
- ✓ Recognizing high-voltage equipment
- ✓ Disposing of components
- ✓ Understanding ergonomics

Before learning how to install computer components such as memory, processors, or a hard drive, you want to make sure that you understand the risks that are involved every time you take the cover off the computer. On the A+ exams, CompTIA will be sure to test your knowledge on how to safely handle a computer's electronic innards and how to safely and legally dispose of broken or obsolete parts.

In this chapter, I discuss topics such as *ESD (ElectroStatic Discharge)* and *EMI (ElectroMagnetic Interference)* as well as concepts such as what high-voltage equipment is and why you should stay away from these components. I also show you how to dispose of computer components and introduce you to ergonomics.

Gathering Tools of the Trade

When plumbers or electricians come to your house to do some work, they bring a toolbox full of the tools they need to fix the vast majority of the problems that might arise while on the job. As an A+ Certified Professional, you need to do the same thing: keep a PC repair toolkit of all the tools, both software and hardware, that you need to do your job.

The following sections discuss the hardware and software tools that you need to perform your job well.

Hardware tools

Of the myriad pieces of hardware you need to repair computers, the most popular is *the screwdriver*. You can purchase a PC repair toolkit (shown in

Figure 3-1) at your local computer store, but, in its simplest form, you should be able to get away with a multihead screwdriver for most jobs.



Figure 3-1:
A typical PC
repair
toolkit.

The following is a list of tools that you may need from time to time to fix most of your computer repair problems:

- ◆ **Screwdrivers:** Most premade toolkits come with a variety of screwdrivers that have a variety of tip shapes and sizes — flat head screwdrivers with a flat tip, Phillips screw drivers with a +-shaped tip, and TORX screwdrivers with a star-shaped tip. I usually use a multi-head screwdriver instead of a mass of different screwdrivers.



Be sure that you don't use a magnetic screwdriver on the computer because you could erase your data!

- ◆ **Tweezers:** PC repair toolkits typically come with a set of tweezers, which are a lot more useful than they might seem. You use tweezers to work with smaller items in the computer, such as jumpers (for more information on jumpers, check out Book II, Chapter 1).
- ◆ **Chip extractor:** A lot of PC repair toolkits come with a chip extractor, which you use to remove a chip from a board. For the most part, you won't need this because most of the chips today are on a card or module instead of being attached directly to the board — and you can remove the entire module by hand.

- ◆ **Chip inserter:** PC repair toolkits also come with a chip inserter, which you use to push a chip into a socket. This tool is just as useful (or just as useless) as a chip extractor.
- ◆ **Canister:** A number of premade PC repair toolkits include an empty canister, which seems really silly at first but becomes very useful. You can use this canister to store additional screws or jumpers. Before throwing away any old computer parts, I usually take the screws and jumpers off and put them in my canister — you never know when you might need them!
- ◆ **Multimeter:** You can use a multimeter (shown in Figure 3-2) to verify that the wall outlet is supplying the correct voltage and that your power connectors to the hard drive and CD-ROM drives are supplying the correct voltage as well. This is a useful tool to have around if you suspect the power supply isn't doing its job.



Figure 3-2: A digital multimeter can be a valuable tool when troubleshooting power problems.

Software tools

Most technicians think of hardware tools when they think of tools of the trade, but software tools are important, too! I give you in-depth information about the software tools you need in your PC repair toolkit in Book VI, Chapter 4, but I want to introduce them here so you get a more complete listing of the tools you will use in the field:

- ◆ **Windows boot disk:** A Windows boot disk contains the core boot files that are needed to boot a Windows 2000/XP/Server 2003 operating system. If those files go missing off the hard drive, you can boot the operating system from the boot disk.
- ◆ **Windows CD:** The Windows CD can help you fix boot problems with Windows when a file goes missing. You can also use it to recover from some basic disk problems by using the Recovery Console.
- ◆ **Live CD:** *Live CD* is the general term used for booting and running the operating system from CD. There are Live CD versions of Windows XP and Linux, and you can use them to help retrieve your data when the operating system will not boot.



These are just a few of the software tools that you can use. You find out more about troubleshooting utilities found in Windows by checking out Book VI, Chapter 4.

Avoiding the Dreaded ESD

In this section, I discuss the dangers of servicing a system without taking precautions as to not damage the computer or yourself. Be sure to take this section very seriously!

ESD means electrostatic discharge (not extrasensory deprivation)

When you walk around on carpet, scuffing your feet, you build up a static charge. You walk over to your best friend and touch his or her hand and zap — an explosive shock! You feel that static shock because you have a higher static charge than your friend. When two objects with dissimilar charges touch, the excess electrons are transferred to the object with the lesser charge. This is when you feel the shock!

Here comes the important part: *That little shock is more than enough to kill a computer chip.* So, had you touched a computer component instead of your friend, you would most likely kill the poor little computer chip.

The scary part about all this is that in order for you to feel and hear the shock, you must have a charge of approximately 3,000 volts, but only about 200 volts are needed to kill a computer chip — just 30 volts can do long-term damage to the chip.

The transferring of the static charge to another object is called *ElectroStatic Discharge (ESD)* and is responsible for damaging or killing computer components, such as computer chips. It is your job as an A+ Certified Professional to make sure that you eliminate ESD by following ESD best practices.

Preventing ESD

You can reduce ESD, but because ESD occurs naturally, you can never completely get rid of the static build-up you carry with you. To reduce your chances of damaging a system when services the computer, follow these best practices:

- ◆ **Use a grounding strap:** One of the most popular solutions for preventing damage to computer components is to wear a grounding strap (shown in Figure 3-3), also known as an ESD wrist strap. The grounding strap goes around your wrist and contains a *resistor* that absorbs any static electricity. The wrist strap then clamps to the metal chassis. The *chassis* is the metal casing of the computer. This gives you constant grounding while servicing the computer, eliminating any static build-up.




Figure 3-3: An antistatic wrist strap, also known as a grounding strap.




If you don't have a grounding strap to connect to the chassis of the computer, be sure to constantly eliminate any static electricity that has built up on yourself by touching the chassis before touching any computer part.

- ◆ **Carpet spray:** You can use a specialized carpet spray on the carpet around your work area to help reduce the static electricity, but it won't eliminate the static entirely. Don't rely on antistatic carpet spray as your *only* source to eliminate static electricity.
- ◆ **Antistatic bags:** When carrying computer components from one place to another, it's a good idea to carry the components in an antistatic bag to protect the components from static electricity. For example, if you are going to carry some memory to a user's workstation, carry the memory in the antistatic bag.
- ◆ **Humidifiers:** You can use humidifiers to place moisture in the air. Your goal is to get the humidity level to about 50 percent. Overly dry environments help create static electricity.

Watching Out for EMI

A number of people confuse ElectroStatic Discharge (ESD) with *ElectroMagnetic Interference (EMI)*. EMI involves computer equipment receiving electrical interference from an outside source, such as another electrical device or a piece of machinery.



Be sure to know the difference between ESD and EMI.

While ESD often causes permanent damage (if you fry a computer chip, it's damaged permanently), EMI is, for the most part, temporary. For example, if you notice that your monitor display is distorted because you are getting interference from an outside source, moving the monitor away from the external source should get the monitor back to normal.

That being said, EMI exposure for a long period of time *can* cause permanent damage to computer components.

Recognizing High-Voltage Equipment

High-voltage equipment is equipment that can cause serious injury if you touch it, and it may even kill you. There are two major computer components that you will not typically service unless you have special training or certification beyond the A+ Certification exam:

- ◆ **The power supply unit**
- ◆ **The monitor**



The danger in servicing the power supply and a cathode ray tube (CRT) based monitor is that both devices use capacitors that hold an electrical charge even after the component is unplugged. The electrical charge is enough to cause severe shock or electrocution. Beware!


High-voltage equipment usually bears a warning label that indicates that you shouldn't open the component and that, if you do, you could be severely hurt. Bottom line, you're playing with fire (all right, electricity) if you open these components. It is best to have them serviced by a certified electrician or someone with special training.

Power supply unit (PSU)

The *Power Supply Unit (PSU)* is responsible for taking *Alternating Current (AC)* from the wall outlet and converting it to a *Direct Current (DC)* that is usable by computer components.

The power supply usually takes 120 volts of alternating current from the wall outlet and supplies the power to computer components by converting the 120 volts AC into 3.3 volts, 5 volts, or 12 volts of direct current.

The power supply contains a charge even after it is unplugged, so servicing a broken power supply isn't recommended. If your power supply is broken, just replace it.



For the exam, remember that a power supply is high-voltage equipment that you should keep your hands and tools away from.

Monitor

The part of the monitor that you need to be aware of is the *Cathode Ray Tube (CRT)*. Creating the images on a CRT monitor takes a lot of power, typically 50,000 volts — which is enough to cause serious damage to yourself, if not kill you. It is important to remember that this charge is held even when the monitor is unplugged from the wall, so, as a general rule, don't open the monitor unless you are a qualified professional (and remember that passing the A+ Certification exam doesn't qualify you to work with CRTs).



As a general rule, you should never open up the monitor — send it to a qualified professional. Due to the seriousness of injury, opening the monitor is not recommended unless you know what you are doing.

For the A+ exams, you need to be familiar with the general steps to discharge a CRT. You would discharge the CRT if you were going to perform some work on the monitor (which this author does not recommend unless you are qualified). If you decide to open a monitor, you will notice a heavy wire running along the side of the CRT — this wire supplies the 50,000 volts.

Be sure *not* to touch this wire! If you are going to work on a monitor, it is recommended that you first discharge the high voltage. To discharge the CRT of a monitor, you need a *shorting probe*. Follow these steps to discharge a monitor's CRT:

- 1. Unplug the monitor and open it up.**
- 2. Clip the wire on the shorting probe to the metal chassis of the computer.**
- 3. Slip the end of the shorting probe under the rubber grommet (looks like a little rubber flap).**

Make sure that you don't touch any other metal parts. You should hear loud popping noises as the high-voltage equipment is discharged.

- 4. When the popping noise has stopped, you may remove the probe.**



Note that the general steps to discharge a CRT are here for A+ exam preparation purposes only. It is not recommended to open a monitor for any reason unless you are qualified.

Disposing of Components

It's important to understand that when you replace an old part from a computer, you have a responsibility to dispose of the old part properly. You are not permitted to dump certain computer components in a landfill because they are hazardous to the environment. For example, the chemicals in a battery could leak into the water supply and contaminate the drinking water. In this section, you find out how to dispose of certain computer components that can be dangerous to the environment.

Federal and local laws regulate the disposal of hazardous materials; if you break them, you could be fined. If you're unsure of how to dispose of a certain component, contact your state or province environment regulator office to find out how to appropriately dispose of it.

You can also contact the U.S. Environmental Protection Agency (EPA) for proper procedures and guidelines on how to safely dispose of computer components. These guidelines are also published on the EPA's Web page at www.epa.gov/epaoswer/osw/hazwaste.htm.

Batteries

Because batteries contain metals such as nickel and cadmium (which are environmentally harmful), federal, state, and local laws prevent people from dumping their batteries in trashcans and throwing them to the curb destined

for a landfill. The federal government passed the Battery Act in 1996 to ensure that the public properly disposes of rechargeable batteries and batteries that contain mercury. You can find information about the Battery Act by checking out www.epa.gov.

Each state has different regulations and rules that deal with the disposal of batteries, and you should be sure that you follow these rules when disposing of old batteries. Contact your state's environmental regulatory office for more information on disposal of batteries and other computer components.

CRTs

CRTs contain harmful components and, therefore, cannot be dumped in landfills. You can donate old monitors that are still in working condition to public schools, churches, and other charities. With any donation, you can get a tax benefit! If the monitor doesn't work or is so badly damaged that it can't be given away, you can contact recycling companies that will take it off your hands and use it for spare parts.

Toner kits and cartridges

A number of vendors offer a small discount if you turn in old toner kits or cartridges when you buy new ones. Even if they don't give a small discount, at least they'll take the old toner off your hands! These vendors typically send the old toner kit or cartridge back to the manufacturer to be recycled.

If the vendor you purchase from doesn't take the used toner kits and cartridges, you should contact the manufacturer and find out how you can dispose of them. If you don't want to contact the manufacturer, you can always check with your state's environmental regulatory office.

Material Safety Data Sheet (MSDS)

A *Material Safety Data Sheet (MSDS)* is a document that contains information about hazardous chemicals that are present in different materials. The MSDS contains the composition of these materials and includes their ingredients so that you will know what makes up a particular material.

The MSDS also contains information on the proper handling of materials and the lethal doses for each material. You can obtain information on MSDS by visiting www.osha.gov and running a search for **MSDS** and your material.

A great Web site to learn more on how to donate or recycle computer parts is www.computerhope.com/disposal.htm.



Exploring Ergonomics

Computers do a lot of great things. They crunch thousands of numbers and process data at phenomenal rates, they perform tedious, repetitive actions without complaining, and they allow a thousand-page book like this one to be written, edited, and proofread without using a single sheet of paper. Unfortunately, working on a computer can cause health problems.

Ergonomics is the study, or science, of designing equipment that reduces discomfort for the user and reduces or prevents *Repetitive Strain Injury (RSI)*.

RSI is an injury to the muscles and tendons in the neck, shoulders, arms, wrists, and/or fingers caused by the poor posture and frequent repetitive movements that accompany long hours in front of a computer. Some common examples of RSI are carpal tunnel syndrome, tenosynovitis, and tendonitis. Here are some recommendations you can follow to help reduce discomfort and RSI:

- ◆ **Maintain good posture:** Be sure to keep good posture when working on the computer (see Figure 3-4). Put your feet flat on the floor; your thighs and back should be at a ninety-degree angle from each other; your forearms should be on the table, parallel to the floor.




Figure 3-4: Having good posture when working on a computer helps prevent injuries to your neck and back.

- ◆ **Use a comfortable keyboard and mouse:** Try to keep the keyboard and mouse at a comfortable position so that you can keep a relaxed posture and not have to stretch too far to do your work. If at all possible, use an ergonomic keyboard, which has a left and a right half set at comfortable angles to help reduce the strain put on your wrists, arms, and fingers.
- ◆ **Monitor at eye level:** Keep the monitor at eye level so that you don't have to bend or stretch your neck to look at it. Keep your monitor two feet away from you because of the rays monitors emit.
- ◆ **Take frequent breaks:** It is so important to take time out in the day to get up and stretch your legs and arms. Staring at the computer screen too long can also hurt your eyes, so give them frequent breaks, as well. It's great for the mind, too!
- ◆ **Care for your eyes:** If you are working at the computer for long periods of time, be sure to blink your eyes frequently and try to give them a break by focusing on a distant object for a few seconds every 15–20 minutes. It has been found that computer users have dry eyes because they don't blink as often as someone not using the computer all day. Have some eye drops handy when working on the computer for long periods of time.
- ◆ **Use a wrist rest:** If you do a lot of typing and mousing on the computer, use a wrist rest for the keyboard and mouse so that you don't put unnecessary strain on your wrist due to bad positioning.
- ◆ **Invest in a good chair:** If you're going to be sitting for a while, you might as well invest in a good chair. Buy a chair that has an adjustable back and adjustable height and that provides lower-back support.

Getting an A+

This chapter highlights a number of safety issues surrounding the servicing and using of computer systems. The following is a quick review of some of the key points to remember from this chapter:

- ◆ ESD can kill computer chips, so make sure that you do your best to protect computer components.
- ◆ Always use a grounding strap when servicing a computer.
- ◆ EMI is electrical interference from an external source.
- ◆ Recycle computer parts or donate them to charity whenever possible.
- ◆ Power supplies and monitors are high-voltage equipment that hold a strong electrical charge even after unplugged, so send them to a qualified technician to be serviced.
- ◆ Be sure to maintain good posture and use a supporting pad for your wrist if you're doing a lot of typing.

Prep Test

- 1** Which of the following is considered high-voltage equipment? (Select all that apply.)
- A Printers
 - B CRTs
 - C Power supply units
 - D Hard drives
- 2** Computer components can be destroyed with as little as:
- A 3000 volts
 - B 200 volts
 - C 500 volts
 - D 900 volts
- 3** ESD stands for:
- A Electrostatic discharge
 - B Electrostationary discharge
 - C Electrostatic disruption
 - D None of the above
- 4** What does EMI stand for?
- A ElectroMagnetic interference
 - B Electric motor interference
 - C Electromechanical interference
 - D Emits magnetic interference
- 5** EPA stands for:
- A Environmental Protection Association
 - B Environmental Protection Agency
 - C Environmental Preservation Agency
 - D Environmental Preservation Association
- 6** RSI stands for:
- A Restive strain injury
 - B Repetitive self-injury
 - C Restrictive system injury
 - D Repetitive strain injury

7 The static shock one commonly feels is typically about how many volts?

- A 3,000
- B 200
- C 10
- D 5

8 You can prevent damage from ESD by:

- A Keeping the air dry
- B Using a multimeter
- C Using a grounding wrist strap
- D Using a static meter

9 A grounding wrist strap contains a:

- A Surge protector
- B Capacitor
- C Voltmeter
- D Resistor

10 The risk for electrostatic discharge is the greatest at:

- A Daytime
- B High humidity
- C Low humidity
- D Nighttime

Answers

- 1 B, C. Examples of high-voltage equipment are CRTs and PSUs. See *"Recognizing High-Voltage Equipment."*
- 2 B. Computer chips and components can be killed with about 200 volts. Review *"ESD means electrostatic discharge (not extrasensory depravation)."*
- 3 A. Electrostatic discharge, or ESD, occurs when two objects carry a charge of different amounts and the charge is transferred from one object to another to balance the charge. Check out *"ESD means electrostatic discharge (not extrasensory depravation)."*
- 4 A. EMI stands for electromagnetic interference and is electrical interference from an external source. Peruse *"Watching Out for EMI."*
- 5 B. The Environmental Protection Agency provides guidelines on how the components should be disposed of. Take a look at *"Disposing of Components."*
- 6 D. When you continuously work with your computer for prolonged periods of time, there is a chance of developing a repetitive strain injury. Peek at *"Exploring Ergonomics."*
- 7 A. The shock we give one another is typically 3,000 volts. Look over *"ESD means electrostatic discharge (not extrasensory depravation)."*
- 8 C. You can prevent damage from ESD by using a grounding wrist strap. Study *"Preventing ESD."*
- 9 D. An ESD wrist strap contains a resistor. Refer to *"Preventing ESD."*
- 10 C. The risk for electrostatic discharge is greatest at low humidity. Examine *"Preventing ESD."*

Chapter 4: An Overview of System Components

Exam Objectives

- ✓ **Introduce the basic components that make up personal computers**
- ✓ **Identify the fundamental principles that apply to personal computers**

This chapter defines basic terms and provides a brief overview of many topics that will be covered in this book. By the time you finish reading this chapter, you should have a good understanding of the major components of the personal computer, from the motherboard to the monitor.

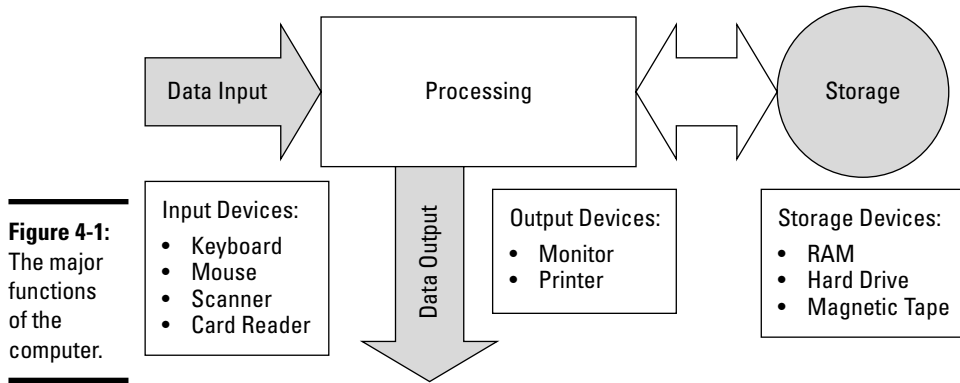
In preparation for the rest of the minibooks, I give you an overview of the major elements that make up the modern computer. As an A+ Certified Professional, you must be aware of these components, be able to service and manage these components, and be able to explain computing concepts to others.

What Is a Computer?

Computers are a major element of our society, and they exist in so many places that we tend to forget that they are there. In simple terms, all computers are made up of four basic functions:

- ◆ **Data Input**
- ◆ **Data Output**
- ◆ **Processing**
- ◆ **Storage**

These functions, along with their related devices, are illustrated in Figure 4-1. In this chapter, I briefly introduce each of these functions and devices, which are covered in depth in Book II, Book III, and Book VIII.



Although the packaging of these functions takes very different shapes, from Personal Digital Assistants (PDAs) and Digital Video Recorders (DVRs) to routers, laptops, and tower computers, they all share these functions at the most basic layer. The CompTIA A+ Certification exam focuses on the common arrangement of devices performing these four functions, which is typically called the *personal computer* and not the more obscure arrangement of devices, such as the automobile engine idle controller. The rest of this chapter guides you through the major components in the “personal computer.”

Looking Inside the Box

Most of the processing and storage devices show up “inside the box.” This section gives you a look at the major elements that you will find inside the system housing, or case.

Processor/CPU

The *processor* (also known as the *Central Processing Unit*, or *CPU*) is the “brains of the organization,” so to speak. It is designed to do very few things, but to do them extremely quickly. The processor performs a limited set of calculations based on requests from the operating system and controls access to system memory. Processor speed is measured in several different ways, including clock cycles, megahertz (MHz), and Millions of Instructions Per Second (MIPS). Any of these measures give you an estimate of the processor’s power.

The speed of early processors ranged from 4 MHz to 8 MHz, while today’s processors have broken the multi-gigahertz (GHz) mark.



You will find processors covered in detail in Book II, Chapter 2.

Storage devices

Storage devices on your computer are responsible for storing data, such as the operating system, applications, and actual output of applications or user data. Depending on the amount and type of data, there are five basic types of devices to work with:

- ◆ **Floppy drives**, including some of the high-capacity formats (such as 120MB Superdisks)
- ◆ **Hard drives**, including some of the removable cartridge drives (such as SyQuest drives)
- ◆ **Optical drives**, including CD-ROM and DVD drives
- ◆ **Magnetic tape drives**, which come in a variety of capacities and are usually used for archiving data
- ◆ **Flash drives**, which store data in a variety of non-volatile memory chips

Any one of these formats enables your computer to store and retrieve data. Each of these storage options is considered to be long-term or permanent storage, but that term is relative as each media format has a limited lifespan.

When dealing with storage devices, there are three major technologies used to connect hard drives to motherboards: *IDE (Integrated Device Electronics)* or *ATA (AT Attachment)*, *SATA (Serial ATA)*, and *SCSI (Small Computer System Interface)*. An ATA controller allows the 40-pin IDE connector to accept a wide variety of devices, while maintaining full backward compatibility with tradition IDE drives or devices. In most cases, IDE and ATA are used interchangeably, while the current systems are technically using ATA technology. There has been a long ongoing battle for speed and performance between IDE (or ATA) and SCSI, but in general, SCSI provides faster and more reliable transportation, while IDE has been the low-cost alternative. SATA is IDE's new cousin, and will help IDE's fight to challenge the benefits of SCSI.



Do not forget that ATA and IDE terms are used interchangeably, but you do not want to confuse these terms with Serial ATA, which uses different technology.

Within a computer a bus is a mechanism that is used to move data between devices of the computer, much like a city bus is used to move people between bus stops in the city. Hard drives are manufactured as ATA, SATA, or SCSI devices, and they are usually connected to the motherboard through the high-speed *PCI (Peripheral Component Interconnect)* bus but being internal devices cannot be easily transported between computers. The USB, IEEE

1394 (or FireWire), and PCMCIA (PC Card) buses provide fast-enough data transfer to allow hard drives to be attached to a system using these technologies, and are designed to be used with external devices. There are many manufactures that have products that allow standard ATA or SATA drives to be connected using these buses. Due to the higher cost and the difference in technology, SCSI drives are not normally used.

The type of drives that you can attach to your computer depends on the types that are supported by your motherboard and I/O cards.



You can investigate what you need to know about hard drives and storage devices for the A+ exam by going to Book II, Chapter 5.

Memory

Memory is a solid state (no moving parts) storage medium. It can take many forms, such as RAM (Random Access Memory), flash RAM, ROM (Read-Only Memory) or EPROM (Erasable Programmable ROM). Access time for RAM is measured in nanoseconds (a billionth of a second). When memory is discussed, it usually refers to *RAM (Random Access Memory)* which is the computer's primary working memory. RAM is a form of storage, although it is usually temporary storage, and many people may not think of it as a storage medium, since it is volatile and loses its information when power is removed, but temporary storage is still storage.



RAM is always temporary, and requires power to retain information. When you put your computer into a hibernation state, the contents of RAM are written to file, and then retrieved when you power up your computer.

An analogy that I often use when talking about RAM is that of a tabletop (RAM) and a filing cabinet (hard dirve or magnetic tape). People will not argue the point that a filing cabinet is a storage area for my papers, but they may dispute that the tabletop is a storage area. When I want to work with my papers, I take them out of the filing cabinet, and place them in temporary storage on my tabletop, where I am able to read the papers and make changes to them. When I am done, I will place the papers back in the permanent storage of the filing cabinet. I know some people who would use their tabletop as a form of permanent storage. Just as with RAM, my tabletop has a limited space, so when it becomes full, I will find that my work gets slower as I have to shuffle papers about. When the OS loads, it loads into RAM; when applications load, they load into RAM; when you open documents, they load into RAM; and when you need to send output to your monitor, the output is loaded into RAM before it appears on your monitor.

In addition to working memory, RAM is used in many areas of your computer. In most cases it is used for caching data or dedicated to specific sub-systems. Most video cards will have dedicated RAM on the video card, while processors and hard drives have special high-speed RAM for caching.



To be able to answer all of the memory questions on the A+ exam, read through Book II, Chapter 3.

System boards

The term *system board* describes any number of circuit boards that make up the internals of your computer, but it is used most often to describe the motherboard. The *motherboard* is the main board in your computer that contains the BIOS chips, RAM, I/O (Input/Output) ports, and CPU. This board maintains the electrical pathways that enable all other components to communicate with each other.

Some service manuals use the term *daughter board* to describe a secondary board which contains motherboard functions such as disk or I/O control. A *daughter board* is a board that contains some of the chips that could have been put on the motherboard but were not — perhaps due to space limitations or for other reasons. It is common to see them used in laptops and other mobile devices.



Motherboard information required for the A+ exam can be perused in Book II, Chapter 1.

Power supply

The *power supply* does exactly what its name suggests: It supplies power to the rest of the components in the computer. The power supply takes 120 or 240 volts (depending on the country you are in) of alternating current from your electrical outlet and converts it to 3.3 volts, 5 volts, and 12 volts of direct current. It contains a number of leads that supply different voltages for different types of devices (such as floppy drives and hard drives).



Book II, Chapter 6 provides you with all of the information you need to answer the power supply questions on the A+ exam.

Adapter cards

Adapter cards allow you to adapt your computer to another role, such as controlling assembly line robots, or to add specific functionality to your computer, such as printing. The most common adapters are now integrated into most motherboards. These components include display, network, keyboard control, mouse control, serial ports, parallel ports, and USB ports.

With the first computers that came out on the market, most of the elements (if not all of them) were added to the motherboard by using the adapter slots, such as ISA, PCI, and AGP (well, ISA anyway, since the other slots did not exist).

Cooling system and fans

As electricity moves through any circuit, heat is generated, which is illustrated by looking at a simple circuit that contains an incandescent light bulb. As electricity moves through the bulb, the filament heats to glowing. Heat is also generated as electricity moves through the integrated circuits that are contained in most of the components of the computer. In addition to these fixed integrated circuits, hard drives contain moving parts that generate heat from friction. These components generate a large amount of heat, which needs to be removed from the critical components to prevent them from failing early in their careers.

To remove heat from computer systems, there has been a steady increase in the use of fans, vents, and other heat dissipation units. Heat sinks with fans are commonly placed on processors and critical chipsets. *Heat sinks* are heat conductive metals (usually aluminum or copper), which have a solid side in contact with the chip that they are protecting, and thin fins on the other side, putting greater surface area with the air, to dissipate or transfer heat to the air. The use of a fan with the heat sink allows more cool air to flow through the fins. Case vents and fans bring cool air into the computer housing and to vent and pull hot air out.

Faster processors produce more heat. As processors have become faster, more methods for pulling the heat away from the processor have been developed.



You will find answers to A+ exam questions related to cooling processors in Book II, Chapter 2, and for overall systems in Book IV, Chapter 1.

Firmware and chipsets

People like to keep their world in a perceived state of order to give themselves a sense of control. This is often seen by the way we classify everything we see or work with into categories. Many of these categories seem very distinct until something comes along to challenge our opinions, and then the waters become murky.

One of these murky areas is the distinction between hardware and software, which at one point were thought to be distinct and separate. *Software* is programming code that is stored on your disk or on some other form of media. *Hardware* refers to the physical components — boards, peripherals, and other equipment — that make up your computer.

Firmware fills in a middle ground between software and hardware, where the distinct line begins to disappear. Firmware is programming code (software) that is contained in or stored on the *IC (Integrated Circuit)* chips (hardware) on your computer. This combination of hardware and software makes up the

BIOS on several different devices, with settings stored in CMOS or flash RAM. This firmware is tied to the function of the IC chips that it is working with; in the case of a network card, firmware would manage network or PXE (Pre-execution Environment) boot functionality of the card.

Most modern motherboards have a series of IC chips and firmware that work together to control the integrated functions of the motherboard. The compatibility of these chips and the code that ties them together is provided by a single supplier, and this group of chips is referred to as a *chipset*. Intel and VIA Technologies produce popular chipsets.

BIOS

BIOS is short for *Basic Input-Output System*. The BIOS is actually software that is stored on a ROM chip on the motherboard. Most systems today use a Flash EPROM (Erasable Programmable ROM) to store the BIOS so that the user can update the programming code in the BIOS.

The BIOS is responsible for controlling or managing low-level but extremely important processes like the POST (Power-On Self-Test), the boot process, and the interaction of components on the motherboard.

Book II, Chapter 4 will provide you with all of the facts that you will need to answer BIOS questions on the A+ exam.



CMOS

CMOS is short for *Complementary Metal-Oxide Semiconductor*, which is the type of manufacturing process that creates most integrated circuits. This development process is used to create the following:

- ◆ **High-density DRAM (Dynamic Random Access Memory)**
- ◆ **High-speed processors**
- ◆ **Low-power devices for mobile use**

The term *complementary* refers to the fact that these chips use negatively and positively charged transistors (which complement each other) to store information. Most RAM chips rely on CMOS technology to store information, but when discussing CMOS, you will probably be referring to the hardware configuration settings that are saved between reboots of your computer. These settings include

- ◆ **Which hard drives and floppy drives are present?**
- ◆ **How much memory is installed?**

- ◆ Is a keyboard required to boot?
- ◆ What type of mouse is installed (PS/2 or serial)?
- ◆ What are the reserved resources (such as IRQ, I/O addresses, and DMA channels)?
- ◆ What is the power-on password and is it required to boot up the system?
- ◆ What are the date and time?
- ◆ Is ACPI (Advance Configuration Power Interface) enabled, and what devices does it apply to?



To be ready for your A+ exam, review the CMOS information, such as common settings, in Book II, Chapter 4.



Remember that BIOS stores programming code, and CMOS stores settings for the BIOS options.

Checking Outside the Box

Now that you have looked at what is inside the box, you will want to see what gets added to the system, outside the box.

Casing and form factors

Part of the outside of the box is the box itself. There are many different form factors for the box, some of which dictate the form factor of the motherboard going into the case. Cases come most often in tower or desktop form factors, but are also found in forms that make them attractive for entertainment units and in extremely small forms for specific uses.

Input and output devices

Computers use many different kinds of input and output devices, which connect to the computer via one of the computer's ports. In the following sections, I give you an overview of the most common input and output devices.

Monitor

The different types of buses that can provide video services include

- ◆ **ISA (Industry Standard Architecture)**, which runs at 8 MHz
- ◆ **PCI (Peripheral Component Interconnect)**, which runs at 33 MHz and 66 MHz

- ◆ **PCIe (Peripheral Component Interconnect Express)**, which is a high-speed serial bus, so its speed is measured differently than the others which are parallel buses. PCIe has between 1 and 16 channels (1x – 16x)
- ◆ **PCI-X (Peripheral Component Interconnect Extended)**, which runs at 133 MHz
- ◆ **AGP (Accelerated Graphics Port)**, which runs at 66 MHz (but can be increased to 8x the base speed or 528 MHz)

You have probably already guessed that the faster the bus speed, the faster your video card is likely to function. The AGP bus was designed specifically for video and is being replaced by PCI Express. In addition to a fast bus speed, video performance and color depth are provided by RAM or video RAM. This RAM is found on the video card itself. Some high-end video cards also have a small processor to handle some of the work of displaying information on your monitor instead of letting the computer's main processor do all the work.

Video cards traditionally allow for color depths that include

- ◆ **4-bit or 16 colors**
- ◆ **8-bit or 256 colors**
- ◆ **16-bit or 65,000 colors**
- ◆ **24-bit or 16 million colors**
- ◆ **32-bit or 4 billion colors**

Standard screen resolutions are (in pixels)

- ◆ **640x480**
- ◆ **800x600**
- ◆ **1024x768**
- ◆ **1152x864**
- ◆ **1280x1024**
- ◆ **1600x1200**

Modern video cards follow the SVGA (Super Video Graphics Array) standard, but all support at least VGA. The VGA standard is output in 16 colors at 640x480.



Review Book III, Chapter 3 to fully understand monitors and video.

Modem

Modem is short for *Modulator/Demodulator*. Modulation refers to the conversion of a digital signal to an analog signal, and demodulation reverses this process. Your computer is digital, while the phone lines that you want to communicate over are analog. In order to allow the digital signal to be passed over the analog lines, you must use a modem.

Modems allow computers to be placed anywhere there are phone lines, and still communicate with each other. Prior to modems, there was no low-cost means of connecting distant computers together; your only choice was an expensive dedicated leased phone line from the Telco. The speeds at which modems operate have been increasing since their invention. They started with transfer rates of 300 bps (bits per second) and have moved up to 115 Kbps.

Today, where modems were traditionally used, they are often replaced by other remote connectivity options, such as ADSL (Asymmetric Digital Subscriber Line), ISDN (Integrated Services Digital Network), and broadband cable. Each of these systems has a device which is called a modem to connect your Ethernet network card to the data network. ADSL and broadband cable send the signals over analog networks and modulate your data signal to previously unused frequencies. ISDN is digital on both sides of the connection, so there is no modulation taking place. The correct term for this device would be a TA (Terminal Adapter) and not a modem.

Standard modems connect to your computer through the serial port and can be synchronous or asynchronous. Most modems that are purchased for a computer are asynchronous.



To prepare yourself for modem questions on the exam, read through Book III, Chapter 2.

Ports

Sailing ports provide a location for ships to load and unload goods from one location to another. On your computer, *ports* act as connection points for cables, which allow for the transfer of data between your computer and another device. Several different types of connectors and cables are used to join devices together. Although the list of devices that communicate through the different types of ports is limitless, some of the basic types of ports and their uses are listed in Table 4-1.

<i>Port</i>	<i>Use</i>
Serial	Connects serial devices such as modems to your computer.
Parallel	Connects parallel devices such as printers to your computer.
Video	Connects a monitor to your computer.
USB	Connects various types of devices to your computer. Devices that used other ports in the past are increasingly being converted to use USB ports. Devices that use this port include printers, modems, mice, keyboards, and scanners.
Keyboard	Connects a keyboard to your computer.
Mouse	Connects a mouse to your computer.



Reading through Book III, Chapter 1 will give you a full appreciation of all of the types of ports, cables, and connectors that are used by computers.

Getting an A+

This chapter provides an overview of the major components found in a computer system, including

- ◆ *System boards*, which contain most of the computer circuitry.
- ◆ *Power supply*, which converts a building's AC to usable DC for the computer.
- ◆ *Processor/CPU*, which efficiently executes instructions for the OS.
- ◆ *Memory*, which holds working data and application code.
- ◆ *Storage devices*, which are long-term or short-term storage areas.
- ◆ *Monitor*, which displays data from the computer.
- ◆ *Modem*, which communicates with other devices.
- ◆ *Firmware*, which is a cross-breed between hardware and software.
- ◆ *BIOS*, which are low-level internal communication routines.
- ◆ *CMOS*, which is a storage area for configuration settings.
- ◆ *Ports*, which are used as connection points for other devices.

Prep Test

- 1 Which of the following components is not usually found on a motherboard?**
 - A BIOS chips
 - B USB drives
 - C Memory
 - D I/O ports

- 2 What type of memory loses its contents when power is turned off on your computer?**
 - A CMOS memory
 - B RAM
 - C ROM
 - D EPROM

- 3 North American power supply converts 120-volt AC to which of the following? (Select all that apply)**
 - A 3.5 volts
 - B 3.3 volts
 - C 6 volts
 - D 12 volts

- 4 What is the primary purpose of the processor?**
 - A To convert digital signals into analog signals
 - B To process signals so that they can be displayed on your monitor
 - C To carry out instructions from the operating system
 - D To convert 16-bit data into 8- or 32-bit data

- 5 What unit is used to measure RAM speed?**
 - A Milliseconds
 - B Gigaseconds
 - C Picoseconds
 - D Nanoseconds

- 6 Which two of the following represent hard drive architectures?**
 - A PCI
 - B IDE
 - C CMOS
 - D SCSI

- 7** What standard do most current monitors follow?
- A VGA
 - B SVGA
 - C CGA
 - D FLAT
- 8** Modems are usually attached to a computer through what type of port?
- A Serial
 - B Sequential
 - C Parallel
 - D Modem
- 9** Firmware is composed of which of the following? (Choose two)
- A Software
 - B Middleware
 - C Hardware
 - D Componentware
- 10** The programs that allow the POST to take place are stored in what?
- A BIOS
 - B CMOS
 - C RAM
 - D POSTOS
- 11** What does CPU stand for?
- A Core Predetermination Utility
 - B Complementary Provider Unit
 - C Central Processing Unit
 - D Co-Primary Uniprocessor
- 12** What does CMOS stand for?
- A Complementary Metal-Oxide Semiconductor
 - B Co-Management of Operating System
 - C Configuration Management and Option Semiconductor
 - D Configuration Memory Option System
- 13** What does SCSI stand for?
- A Serial Component System Interface
 - B Small Component Serial Interface
 - C Serial Computing Storage Interface
 - D Small Computer System Interface

Answers


- 1 B. The motherboard will include I/O ports or connectors for a variety of devices, but will not typically include the devices themselves. It will include the ATA controller, and a connector to attach the ATA drive, but not the drive itself. The same would be true of USB drives; the motherboard will include the controller and connector for the USB bus, but not the actual device. See *"System boards."*
- 2 B. RAM memory is volatile or non-permanent as it loses its contents when power is turned off on the computer. *Review "Memory."*
- 3 B, D. Power supplies have leads that supply 3.3 volts, 5 volts, or 12 volts. *Check out "Power supply."*
- 4 C. Processors carry out instructions that they receive from the operating system. *Peruse "Processor/CPU."*
- 5 D. The normal measurement for RAM is nanoseconds, or billionths of a second. *Take a look at "Memory."*
- 6 B, D. Hard drive architectures used with computers are IDE (or ATA), Serial ATA (SATA), and SCSI. *Peek at "Storage devices."*
- 7 B. Super Video Graphics Array is the current graphics standard that monitors follow. *Look over "Monitor."*
- 8 A. Modems are serial devices, and as such are attached to the serial port. *Study "Modem."*
- 9 A, C. Firmware consists of programming code or software that is contained within a hardware component such as a ROM chip. *Refer to "Firmware and chipsets."*
- 10 A. BIOS contains the programs that allow the POST to take place. The BIOS is stored on a ROM or an EPROM. *See "BIOS."*
- 11 C. Central processing unit. *Review "Processor/CPU."*
- 12 A. Complementary metal-oxide semiconductor. *Check out "CMOS."*
- 13 D. Small computer system interface. *Peruse "Storage devices."*

Book II

Inside the Box

The 5th Wave

By Rich Tennant



Contents at a Glance

Chapter 1: Knowing Your Motherboard	65
Chapter 2: Picking Your Processor	115
Chapter 3: What to Remember about Memory	151
Chapter 4: Telling Your BIOS from Your CMOS	179
Chapter 5: Working with Storage	207
Chapter 6: Working with Power	271

Chapter 1: Knowing Your Motherboard

Exam Objectives

- ✓ Distinguishing motherboard components
- ✓ Recognizing types of motherboards
- ✓ Identifying bus architectures

One of the major replaceable components in your computer is the system board, also known as the motherboard. The *motherboard* is the big green board (that may not be a technical description, but I think that looking inside your system will demonstrate that it is an accurate one) connected to the computer case — it is the motherboard that holds your RAM, processor, and a number of other components in place.

This motherboard is the glue that connects all the other PC components together. For example, you can see how the hard drive is connected to the motherboard by following the IDE ribbon cable from your hard drive to the motherboard. If you do the same with the IDE cable that connects to the floppy drive, you can see that the floppy drive also connects to the motherboard. The memory sockets and the processor socket are likewise located on the motherboard.

All the components that work together to make the computer functional connect to the motherboard. If you take a close look at the motherboard, you can see wires embedded on the board that form little pathways that span the system. Think of these wires as the highway system that data signals use to travel from one location to another.

In this chapter, I introduce you to the different types of components found on the motherboard. After identifying the motherboard components, you find out about the different types of motherboards. Finally, you explore what an expansion bus is and discover the different bus architectures.

Finding Out What's on a Motherboard

When you look at the motherboard inside your computer, you notice that a number of different items connect to it. The memory sockets, the CPU socket, and the BIOS chip are all located on the motherboard. In this section, we identify the different components that are interconnected via the motherboard.



Remember that the terms *system board* and *motherboard* are interchangeable.

Processor

One of the easiest items to identify on the motherboard is the processor, also known as the *central processing unit* (CPU). The processor is usually the largest chip on the motherboard and is one of the few chips with a heat sink or fan on top of it, as shown in Figure 1-1.

Processor




Figure 1-1: Identifying the processor on the motherboard.

The motherboard has a socket that the processor is inserted into. Today, this socket is implemented as a *zero insertion force* (ZIF) socket, which means that the processor chip can be removed or added to the socket with very little effort. ZIF sockets (shown in Figure 1-2) typically have a lever that you pull to pop the processor out of the socket.

When the Pentium II processor was developed, Intel used a different type of packaging, known as the *Single Edge Contact* (SEC). Motherboards had to implement a different type of “socket,” known as *slot 1*, to hold this processor. The cartridge would drop into the slot, as shown in Figure 1-3. For more information on processors and sockets, check out Book II, Chapter 2.

Socket 7 ZIF Socket




Figure 1-2: Looking at a ZIF socket located on the motherboard.

Book II
Chapter 1

Knowing Your
Motherboard



Figure 1-3: Looking at a Pentium II using the SEC packaging.



Remember that classic Pentium chips are inserted into socket 5 or socket 7, whereas Pentium II processors are inserted into slot 1. With newer Pentium processors, such as Pentium III and Pentium 4, Intel has moved away from the SEC. The Pentium III is placed in Socket 370 while the Pentium 4 is placed in Socket 423 or Socket 478.

SIMM/DIMM sockets

When you look at a motherboard, one of the first items that should stand out is the processor; the next things you will usually notice are the memory slots that are used to install RAM into the system.

There are typically two types of sockets to install memory: *Single Inline Memory Module (SIMM)* sockets and *Dual Inline Memory Module (DIMM)* sockets. Original Pentium systems typically have either four 72-pin SIMM sockets or two 168-pin DIMM sockets to install memory, while newer motherboards today have up to four DIMM sockets and no SIMM sockets. There are no rules as to how many SIMM or DIMM sockets a motherboard manufacturer may use, as you can see with Figure 1-4. Figure 1-4 shows a motherboard with four 72-pin SIMM sockets *and* two DIMM sockets used to hold memory. SIMMs have been phased out and are only available on older motherboards.




Figure 1-4: Identifying SIMM and DIMM memory slots on the motherboard.

When installing SIMMs in Pentium motherboards, you have to install them in pairs, but when installing DIMMs, you can install them individually. The reason for the difference is that when installing memory, you must fill a memory bank, which is the size of the processor's data path. That is, if you install 72-pin (32-bit) SIMMs onto a Pentium (64-bit) motherboard, then you have to install two modules to fill the 64-bit data path of the processor. DIMMs are 64-bit memory modules, the same number of bits as the data path of the CPU, which is why you're able to install only one at a time. For more information on memory banks and installing memory, check out Book II, Chapter 3.

Cache memory

Cache memory increases performance by storing frequently used program code or data that can be later accessed by the processor. Cache memory is much faster memory than normal RAM and, as a result, is more expensive. The system stores data accessed from RAM in cache memory when the data is accessed the first time, making subsequent requests to the same data faster because the data is accessed from cache (which is faster than RAM) for subsequent calls.

All processors today have integrated cache memory, which is known as *level 1 cache*. *Integrated cache* is cache memory that is built into the processor, while nonintegrated cache — known as *external cache* — is built outside the processor, typically on the motherboard. The types of cache are as follows:

- ◆ **L1 (level-1) cache:** Cache that is integrated within the processor.
- ◆ **L2 (level-2) cache:** Cache that is located outside the processor, usually on the motherboard.


Older motherboards implemented cache memory as rows of DIP (dual inline package) chips placed directly on the motherboard. This area was sometimes even labeled “cache.” Unfortunately, you can't expect a motherboard to be well-labeled; if you find labels (in English), consider it an added bonus! For more information on chip packages, check out Book II, Chapter 2.

Other systems have implemented the cache as a memory module, so you may see an empty slot on the motherboard that looks like a SIMM slot, but it will really hold a cache module. A lot of times, this will be labeled as “cache” on the motherboard. Figure 1-5 shows L2 cache on an older motherboard.

Remember, L2 cache is usually located on the motherboard near the processor. That way, data travels over a shorter distance from cache to processor — increasing overall system performance. Also, today's processors implement both L1 and L2 cache in the casing of the processor. For more information on cache memory, refer to Book II, Chapter 3.



Figure 1-5: Looking at L2 cache located on the motherboard.




Motherboard chipset

Each hardware component in the system has circuitry that is responsible for managing a specific hardware part. This circuitry is known as the *controller* for that specific piece of hardware. For example, access to memory is controlled by the *memory controller*, the hard disk is managed by the *hard disk controller*, and the keyboard is managed by the *keyboard controller*.

The combination of computer chips that hold the logic for these controllers is known as the *motherboard chipset*. Together, the computer chips make up the chipset control communication from the CPU to each of the hardware devices in the system.


Two chips that make up a big part of a motherboard's chipset are the North Bridge and the South Bridge. The *North Bridge* chip is responsible for communication from the CPU to memory and the *advanced graphics port* (AGP) device (more on AGP later in this chapter). The *South Bridge* chip is responsible for communication between the CPU and other devices, such as PCI, ISA, and IDE devices. These two chips contain the bulk of the logic that allows a CPU to communicate with other hardware. Figure 1-6 displays the relationship between the processor and the North Bridge and South Bridge chips.

Figure 1-6: Looking at the relationship between the CPU and the motherboard chipset.



Locating the North Bridge and South Bridge chips on a motherboard can sometimes be challenging. The North Bridge chip is typically the second largest chip (after the processor) and typically contains a heat sink or fan on top of the chip to keep it cool. The North Bridge chip is typically located between the processor and the AGP slot, while the South Bridge is normally located farther from the processor — usually beside the PCI slots, as shown in Figure 1-7. Notice in the figure that the North Bridge chip has the words AOPEN on it, while the South Bridge is the chip above the PCI slots.

Figure 1-7: Identifying the North Bridge and South Bridge chips on an AOPEN motherboard.



BIOS chip

The *Basic Input-Output System (BIOS)* is the low-level program code that allows all the system devices to communicate with one another. This low-level program code is stored in the BIOS chip on the motherboard.

Locating the BIOS chip on the motherboard is easy; it is usually rectangular and generally features a label with the manufacturer's name and the year the chip was manufactured. Some of the popular manufacturers are AMI, AWARD, and Phoenix.

The BIOS chip is a *Read-Only Memory (ROM)* chip, which means that you can read information from the chip, but you can't write to the chip under normal circumstances. Today's implementations of BIOS chips are *EEPROM (Electrically Erasable Programmable ROM)*, which means that you can get special software from the manufacturer of the BIOS to write to the chip.

Why would you want to erase the BIOS? Suppose, for example, that your BIOS is programmed to support a hard disk up to 2GB in size, but that you want to install a new, larger hard disk instead. What can you do about it? You can contact the BIOS manufacturer and get an update for your BIOS chip, which is usually a software program (in the past, you generally had to install a new chip). Running the software program writes new instructions to the BIOS to make it aware that there are hard disks bigger than 2GB and provides instructions for dealing with them. But before new instructions can be written, the old instructions need to be erased.

The BIOS chip also contains code that controls the boot process for your system. It contains code that will perform a *power-on self-test (POST)*, which means that the computer goes through a number of tests, checking itself out and making sure that it is okay. After it has made it past the POST, the BIOS then locates a bootable partition and calls on the master boot record, which loads an operating system. Figure 1-8 shows a BIOS chip on a motherboard. For more information on the system BIOS, refer to Book II, Chapter 4.

Battery

The computer keeps track of its inventory in what is known as the *Complementary Metal-Oxide Semiconductor (CMOS)*. CMOS holds a listing of system components, such as the size of the hard disk, the amount of RAM, and the resources (IRQs and I/O addresses) used by the serial and parallel ports.




Figure 1-8: Looking at a BIOS chip located on the motherboard.

This inventory list is stored in what is known as CMOS RAM, which is a bit of a problem because RAM loses its content when the power is shut off. You don't want the computer to forget that it has a hard disk or forget how much RAM it has installed. To prevent this sort of problem, a small watch-like battery on the motherboard maintains enough energy that CMOS RAM doesn't lose its charge. If CMOS RAM loses its charge, it results in the CMOS content being lost. Figure 1-9 identifies a battery on the motherboard. For more information on CMOS, check out Book II, Chapter 4.

Expansion slots

Most motherboards have one or more *expansion slots*, which serve the purpose of adding functionality to the computer. For example, assume that your computer doesn't have sound capability — you can install a sound card into the expansion slot to add that capability to your system.

CMOS Battery




Figure 1-9: Identifying the battery on the motherboard that is used to maintain a charge to CMOS RAM.

Expansion slots come in different varieties, and it is extremely important to understand the benefits of each type. We discuss these issues later in the chapter, in the section titled “Understanding Bus Architectures.” For now, I just want you to be able to identify the expansion slots on the motherboard.

If you look at the motherboard, you can see a number of expansion slots. There are probably some white, narrow PCI slots on the board, as well as a tan-colored AGP slot (used for video cards). You may also see some larger black slots; these are ISA slots used by older devices. Most motherboards today do not have ISA slots, or may only have one. Figure 1-10 displays ISA, PCI, and AGP expansion slots used to add expansion cards to the system. For more information on expansion slots, refer to the “Understanding Bus Architectures” section, later in this chapter.

Ports and connectors

There are a number of *ports* on the back of the motherboard that connect the keyboard, mouse, printer, and other devices to the system. This section identifies those ports. Figure 1-11 displays a number of built-in *input/output* (I/O) ports on the back of an ATX motherboard.




Figure 1-10: Identifying expansion slots such as AGP, PCI, and ISA on a motherboard.



Figure 1-11: Looking at the built-in ports on the back of an ATX motherboard.

Serial ports


Most motherboards have serial ports integrated directly into the board. The *serial ports* are also known as *communications (COM) ports*. The reason that they are called *serial ports* is because they send data in a series — a single bit at a time. If eight bits of data are being delivered to a device connected to the COM ports, then the system sends the eight bits of data, one bit at a time, in single file. Typically, there are two COM ports — COM1 and COM2 — on each system.

The official standard that governs serial communication is known as *RS-232*, and you may see serial ports referred to as *RS-232 ports*.



You usually connect an external modem or a serial mouse to a serial port. Each of these devices is used for communication; a modem allows your computer to talk to another computer across phone lines, while a serial mouse allows you to communicate with the system. Figure 1-12 shows two serial ports connected to a motherboard.

Figure 1-12: Identifying the integrated ports on the back of an ATX motherboard.



Serial ports on the back of the motherboard are one of two types:

- ◆ **DB9-male** is a serial port with 9 pins.
- ◆ **DB25-male** is a serial port with 25 pins.

Parallel port

Another type of connector that you will have on the back of the motherboard is the parallel port. The *parallel port* is also known as the *printer port*, or *LPT1*. The parallel port gets its name by being able to send information eight bits at a time. Whereas serial ports send only one bit at a time in single file, parallel ports can send eight bits in one operation — side-by-side rather than single file. Refer to Figure 1-12 to see a parallel port connected to a motherboard.

The parallel port, which is known as DB25-female, has 25 pins and is located on the back of the motherboard. Looking back at Figure 1-12, you can see the parallel port located above the two serial ports.



It is important not to confuse the serial port with the parallel port. The serial port is a male port (meaning that the port has a number of pins in it), whereas the parallel port is a female port (meaning that it does not contain the pins but the pin holes).

You connect the parallel port to a printer by using a parallel cable that has a different type of connector at each end. On one end of the cable is a DB25 connector that attaches to the parallel port on the back of the computer. The other end of the cable (the end that connects to the printer) has a 36-pin Centronics connector.




Remember, a standard printer cable has a different type of connector on each end. One end has a DB25-male connector with 25 pins, while the other end has a 36-pin Centronics connector.

Video adapter

In the past, a motherboard came with a built-in *video adapter*, sometimes called a *video card* or *video controller*. The video adapter is responsible for converting digital data from the processor and preparing the information to be displayed on the screen. Figure 1-13 displays a video adapter port — you can identify it by the three rows of five pins. The video port is a 15-pin female port.

Figure 1-13: Identifying the video adapter port on the system.



Many systems today use the ATX motherboard form factor and, as a result, have an AGP slot to hold the video adapter. This means that the video adapter is not integrated into the motherboard like it was in the past.

Figure 1-14 shows how information flows from the computer system to the monitor. The following steps refer to the numbers in Figure 1-14.

1. The video adapter is responsible for receiving digital data from the processor, which instructs the video adapter on how the images are to be drawn on the screen.
2. The video adapter stores the information about drawing the images in its memory and starts converting the information into analog data that the monitor can understand.
3. The data is sent in analog format from the video adapter to the monitor.




Figure 1-14: Looking at how information flows from the processor to the display.

Keyboard/mouse connectors

The mouse and keyboard connectors on motherboards today are most likely PS/2 style connectors or USB connectors. Let's focus on PS/2 connectors for a now. A PS/2 connector is a small circular six-pin connector. In Figure 1-15, you can see the keyboard and mouse connectors on the left side of the diagram.

Figure 1-15: Identifying the PS/2 connections used for a keyboard and mouse.



Keyboard connector

Mouse connector

Older motherboards may have a *DIN keyboard connector*, also known as an *AT connector*, which you can see on AT and Baby AT motherboards. These systems did not have any other ports on the back of the system, so you would need to insert an I/O card for other ports (such as serial and parallel ports).

Sound

Most motherboards today have built-in sound capabilities, allowing you to connect speakers and a microphone to the computer. Figure 1-16 shows the integrated sound ports on a motherboard. There are three different ports on the integrated sound card:

- ◆ **Line-in:** The line-in port is normally blue and allows you to connect many audio sources to the system.
- ◆ **Line-out:** The speaker port is normally green and allows you to connect speakers to the computer.
- ◆ **Microphone:** The MIC-in port is red and allows you to connect a microphone to the system for recording purposes.

Network interface card and modem

A number of systems today have built-in network support via an integrated *network interface card* (NIC), or network card for short. These systems may have a built-in modem as well. The built-in network card has an RJ45 port on the back of the system that looks like an oversized telephone jack, as shown in Figure 1-17.

Figure 1-16: Identifying the sound ports on the system.





Figure 1-17: Identifying the RJ45 port.



USB ports

Universal Serial Bus (USB) is a high-speed serial technology that transfers data at 12 Mbps (USB 1.0) and 480 Mbps (USB 2.0). One of the major benefits of USB is the fact that all USB devices use the same type of connector, so you won't have to guess which ports to connect the mouse, keyboard, or scanner to. If they are all USB devices, they connect to the same type of port!



The newest USB standard, called USB 2.0, has a transfer rate of 480 Mbps, which is much faster than the USB 1.0 standard. Be sure that you also have USB 2.0 drivers installed to leverage the performance benefits of your USB 2.0 devices.

USB devices also support *daisy chaining*. For example, you can connect Device A to the back of the computer and then connect Device B to Device A, and so on. You can connect up to 127 devices to a system using USB. Figure 1-18 identifies the USB ports on the back of an ATX motherboard.

Figure 1-18: Identifying the USB ports on the back of the motherboard.



USB ports

A USB device that connects to the computer and then has other devices connected to it is considered a *hub device*. If you don't have a USB device that can act as a hub device, you can purchase a specific USB hub that allows you to chain four or more other devices off of it. With a USB hub, you can easily increase the number of USB ports your system has by connecting USB devices to the hub and connecting the hub to the back of the computer. Figure 1-19 shows a USB hub.

Figure 1-19: A USB hub with four ports.



FireWire (IEEE 1394)

When USB 1.0 was introduced, it ran at 12 Mbps. This was a fairly good speed for most types of devices but was a little too slow when it came to multimedia devices, such as digital video cameras. Typically, these types of devices use a *FireWire* connection, which has a transfer rate of up to 400 Mbps and supports 63 devices in a chain. This is a huge jump compared to the USB 1.0 standard. The official standard that defines FireWire is known as the IEEE 1394 — be sure to remember that for the exam!

Just as USB had a second version with a faster transfer rate so does FireWire. The second version of FireWire is defined as the IEEE 1394b standard and transfers data at 800 Mbps! This second version of FireWire is also known as FireWire 800. Figure 1-20 shows a digital video camera being plugged into a FireWire port.



Figure 1-20:
A digital video camera being connected to a system by the FireWire port.



For the exam, remember that the original version of FireWire runs at 400 Mbps and is known as IEEE 1394. The second version of FireWire, FireWire 800, is also known as IEEE 1394b and runs at 800 Mbps. Also, FireWire supports 63 devices in a daisy chain.

For more information on common ports and connectors such as keyboard, mouse, serial, parallel, USB, and FireWire refer to Book III, Chapter 1.

Power connectors

All the devices connected to the motherboard need to get power from somewhere, so the power supply is connected to the motherboard, which supplies power to the board and its components. The following sections discuss power connectors on older and newer motherboards.

Older motherboard power connectors

Figure 1-21 shows power connectors on an older motherboard. There are power cables coming from the power supply to connect to the motherboard with very unique connectors on the end. These power connectors coming from the power supply that connect to the motherboard may be labeled as P1 and P2, or on some systems, P8 and P9.




Figure 1-21: Looking at the power connectors on an older motherboard.

You have to be extremely careful to make sure that the connectors on the cable coming from the power supply to the motherboard are inserted properly, or you could damage the motherboard. Often, these connectors are *keyed* (meaning that they can go in only one way) so that you cannot put both of the connectors in the wrong way. These older power connectors supplied power in 5 volts and 12 volts.

ATX power connectors

Newer ATX motherboards use a different power connector than the one shown in the preceding section. The ATX power connector supplies 3.3 volts, 5 volts, and 12 volts. The ATX power connector, shown in Figure 1-22, is typically labeled as P1.

ATX power connector



Figure 1-22:
The ATX power connector on an ATX motherboard.

Some systems, like ones that use the Pentium 4 boards, use an additional power connector, known as the P4 connector, which supplies an additional 12 volts to the ATX board. Figure 1-23 displays the P4 power connector.




Figure 1-23:
The P4 power connector on an ATX motherboard.



For more information on power supplies and their connectors, check out Book II, Chapter 6.

Drive connectors

You need to be able to identify the different types of connectors that link hard drives to your system. As you may already be aware, the hard drives are used to store information permanently on the computer but in order to access that information the drives have a physical connection to the system via the motherboard.

There are four major types of drives in systems today, *IDE (Integrated Drive Electronics)* drives, *SATA (Serial Advanced Technology Attachment)* drives, SCSI drives, and floppy drives. Each type of drive has its own type of connection on the motherboard. Before you purchase a hard disk to add to the system, you need to be aware of what types of drives your motherboard supports.

IDE connections

IDE drives have been around since the 1980s, and although the technology has improved from a performance perspective, IDE drives connect to the system in the same way they always have. If your motherboard supports IDE, you will have two IDE connectors that are made up of 40 pins each, as shown in Figure 1-24.




Figure 1-24: IDE connectors on the motherboard.

You will connect the drive to the connector on the motherboard by using a 40-wire or 80-wire *IDE ribbon cable*. This ribbon cable typically has two connectors on it — one end connects to the drive, while the other end connects to the motherboard. You can also find IDE ribbon cables with three connectors that allow you to connect two drives to each IDE connector on the motherboard. This means that you can have up to four IDE devices on a system. Figure 1-25 shows an IDE ribbon cable connector.



Figure 1-25:
A 40-wire
IDE ribbon
cable.

When connecting the IDE ribbon cable to the drive and motherboard, the colored wire on the ribbon cable connects to pin 1 on the connector. This is known as the *pin-1 rule*. Pin 1 is normally labeled on the motherboard and drive. If it isn't labeled, see whether the manufacturer has labeled pin 40 — if so, pin 1 is at the other end!



IDE controllers is a popular term used in the computer industry for the IDE connectors. Although in theory these are not controllers, it is a term used in the industry to describe the IDE connections on the motherboard. The *actual* IDE controller is the circuitry located on the circuit board on the drive itself — it is responsible for controlling the flow of information to and from the drive.

SATA connections

Limitations of the IDE architecture have kept its data transfer rate around 150 MBps. As drives become more powerful, a new standard is needed. The first new standard to replace IDE is known as *SATA* and is now becoming popular in desktop computers. SATA can reach transfer rates of up to 600 MBps! This is quite a bit (450 megabytes, to be exact) faster than the 150 MBps currently offered by high-end IDE drives.

SATA uses its own unique four-wire cable to connect to the motherboard. Figure 1-26 shows a SATA cable connected to the SATA connector on the motherboard. Notice that the cable is quite a bit thinner than the IDE ribbon cable; this allows for better airflow in the system and improves overall temperature control of the computer.




Figure 1-26:
A SATA
cable
connected
to the
mother-
board.



Unlike IDE drives, you cannot connect more than one SATA drive to a connector. For example, if your motherboard has two SATA connectors then you can only connect two SATA drives to the system unless you purchase a SATA card that has additional connectors.

SCSI controller

Some high-end machines, particularly those designed for use as servers, may have a controller on the motherboard with 50 pins on it. This is the footprint of a *SCSI (Small Computer System Interface) controller*. Because SCSI devices outperform IDE devices, SCSI controllers are extremely popular for servers (which have greater hard disk access and storage needs than regular desktop computers). To connect a SCSI drive to the 50-pin SCSI connector on the system, you use a 50-wire ribbon cable.



Remember, IDE uses a 40-pin connector that a 40/80-wire ribbon cable connects to, and an internal SCSI connector has 50 pins that connect to a 50-wire ribbon cable.



Lab 1-1 and Lab 1-2 will help you identify the major motherboard components on the motherboard. Lab 1-1 and Lab 1-2 can be found in the Labs.pdf file in the Author directory of the CD-ROM.

Floppy disk connectors

Located very close to the IDE connectors on the motherboard, you should see a smaller *floppy drive connector* which contains 34 pins instead of the 40 pins found with the hard drive IDE connectors. The floppy drive connector on the motherboard is used to connect the floppy drive to the motherboard using a 34-wire ribbon cable.

When connecting the floppy drive to the system, you will notice that the wires on one end of the ribbon cable are twisted. This twisted end must be connected to the floppy drive. The opposite, untwisted end connects to the motherboard. Also note that one wire of ribbon cable is colored, usually red, which indicates wire one. Like the IDE drives, you need to connect wire one to pin 1 on the motherboard and on the floppy drive.



To find out more information about IDE, SATA, SCSI, and floppy drives, check out Book II, Chapter 5.

Jumpers and DIP switches

A jumper is a set of pins that have a plastic cap enclosed over them to create an electrical connection. The plastic cap contains a piece of metal that makes contact with the pins and creates the electrical circuit. The circuit that is created enables a feature on the motherboard. Most motherboards (and expansion cards) use *jumpers* to implement different settings. Figure 1-27 displays a jumper on an expansion card.




Figure 1-27:
Identifying a
jumper.

Notice in the figure that the jumper has three sets of pins that the cap may be placed over. The idea behind the three sets of pins is that each set of pins would enable a different setting. For example, looking back at Figure 1-27, the three different sets of pins may be used to assign three different IRQs to the card. You choose which IRQ is assigned to the card by setting the jumper over a set of pins. Keep in mind that you no longer assign IRQs with jumpers, it was something that was done years ago (for more information on IRQs, check out Book III, Chapter 4).

Today you will find jumpers on motherboards, hard drives, CDROM drives, and DVD drives. There are many different features that can be enabled or disabled on a motherboard using jumpers. For example, there usually is a jumper on the motherboard that is used to clear the CMOS password of a system, to change the voltage supplied to the processor socket, or to change the speed of the motherboard. In order to know what jumper to set you need to check the documentation for the motherboard.

Another popular component of a motherboard or expansion cards in the past that was used to enable or disable different features is the *dual inline package (DIP) switch*. A DIP switch (seen in Figure 1-28) is a set of switches that can be turned on or turned off to enable functionality on the board. In order to know what to set for on/off combinations you would need to consult the documentation for the board.




Figure 1-28:
Looking at a
DIP switch.

Identifying the Types of Motherboards

Now that you understand some of the major components of the motherboard (system board), it is important to mention the different motherboard form factors. A *motherboard form factor* just describes the dimensions of the motherboard and the layout of the motherboard components.

It is important to understand the different motherboard form factors because you can't just take *any* motherboard and place it in a computer case. You must put a full AT motherboard in a full AT case, a Baby AT board in a Baby AT case, and an ATX board in an ATX case. Figure 1-29 shows the three major types of motherboards and gives you an idea of size and shape differences between the three types.




Figure 1-29: Looking at different motherboard form factors.

Full AT

The first type of motherboard that we want to talk about is the *full AT* motherboard. The full AT motherboard is 12 inches wide and 11 inches long and is easily recognized by the fact that it only has a keyboard connector on the back of the motherboard — it contains no other I/O ports.

The full AT suffers from a problem with accessing some of the items on the motherboard because the drive bays hang over the motherboard. This situation makes installation and troubleshooting of the components on the motherboard very difficult.

Another problem with the layout of the full AT board is that the expansion cards, once inserted into the systems, cover the processor. This situation leads to cooling problems because ventilation is insufficient to keep the chip from overheating. Figure 1-30 displays a full AT motherboard being installed in a full AT case.




Figure 1-30: Looking at a full AT motherboard.

Baby AT

The *Baby AT* motherboard form factor had been one of the most popular motherboard types until recent years. The *Baby AT* board is 8.5 inches wide and 10 inches long. This motherboard can be easily recognized because it usually has a DIN keyboard connector in the top-right corner of the board. This keyboard connector is the only I/O connector on the back of the motherboard.

The *Baby AT* board is about two-thirds the size of the full AT board and typically incorporates a socket 7 ZIF (zero insertion force) slot for classic Pentium processors. The *Baby AT* board usually has a mixture of ISA/EISA and PCI slots located on the motherboard and includes a Plug and Play BIOS. Figure 1-31 shows a *Baby AT* motherboard and identifies the popular components.

Take a minute to consider some of the key components on the *Baby AT* motherboard. You can see the socket 7 ZIF slot at the bottom of the motherboard where the processor is to be installed. Also notice the SIMM and DIMM sockets on the right side of the motherboard, which house the system memory. To the left of the SIMM and DIMM slots, you can see the primary

and secondary EIDE connectors (sometimes called controllers) for connecting the hard drives to the board. To the left of the EIDE controllers, notice the types of expansion slots that are used: There are four PCI slots and three EISA slots. Above the PCI slots, you can also see a silver circle, which is the CMOS battery.




Figure 1-31: Identifying components on a Baby AT motherboard.

LPX/NLX

In an effort to allow computers to take up much less space, a slimline desktop system was designed with a smaller motherboard. After the era of the Baby AT came the *LPX* (*low profile extended*), which was then replaced by *NLX* (*New Low-profile eXtended*) motherboard. Both motherboard types served the same purpose — to create low-profile computers.

The NLX motherboard is identifiable by the I/O ports along the back of the motherboard. This motherboard is unlike the full and Baby AT because they incorporated only the keyboard connector. The NLX provides a keyboard and mouse connector, serial and parallel ports, and a video connector.

The NLX form factor is 9 inches wide by 13.6 inches long and uses a riser card to house the bus architectures. The riser card typically connects to the side of the motherboard and is then secured along the side of the case. Figure 1-32 shows an NLX motherboard with a riser card. (We cover bus architectures in the section, “Understanding Bus Architectures,” later in this chapter.)




Figure 1-32:
An NLX
form factor
mother-
board.

ATX

In 1995, Intel wanted a motherboard that would support the Pentium II processor and the new AGP slot, so the ATX form factor was built (shown in Figure 1-33). The ATX board is 7.5 inches wide and 12 inches long and has most of the I/O ports integrated directly into the board, including USB ports.




Figure 1-33:
The position
of the I/O
ports on an
ATX mother-
board.



Remember, the ATX motherboard incorporates the I/O ports and includes an AGP slot for high-performance video cards. Figure 1-33 displays the ports on the back of the ATX motherboard — they are clustered in the left corner of the board and do not spread across the length of the board like they do with the NLX form factor.

The ATX board introduced a 100 MHz system bus and has been increased to speeds of 533 MHz and more. The ATX motherboard has one AGP slot for the video card, which means that the built-in I/O ports on the back of the board do not have a built-in video card like the NLX. The ATX board also has *soft power support*, which allows software developers to create software that controls the startup and shutdown of the system.

The ATX form factor rotated the Baby AT components by 90 degrees so that any cards inserted into the bus architectures would not cover the processor and prevent proper cooling. Figure 1-34 shows an ATX motherboard.

Figure 1-34 also highlights some of the common components on the ATX board. Notice, for instance, slot 1, where a Pentium II chip can be inserted. Newer versions of the ATX motherboard use a ZIF socket to house the processor. Notice also, in the top-right corner, the BIOS chip with a white label on top of it. At the top of the figure, you can identify the EISA and PCI slots, and located in the center of the board is an AGP slot. The hard drive controllers are located on the left side beside the three slots that hold the DIMM memory.





Figure 1-34: The ATX motherboard is very popular in today's systems.

MicroATX and FlexATX

Smaller versions of the ATX motherboard, known as MicroATX and FlexATX, have been developed. The *MicroATX* motherboard form factor is 9.6 inches by 9.6 inches and can fit in a MicroATX case or full ATX case. The *FlexATX* is smaller than the MicroATX (9 inches by 7.5 inches) and fits in an ATX and MicroATX case. FlexATX is not as popular because the size of the motherboard limits how much you can expand on the system. Figure 1-35 shows a MicroATX board.

Figure 1-35: Comparing the size of the ATX motherboard (left) with the MicroATX motherboard (right).



The important point to make here is that when you purchase a motherboard, you must ensure that the motherboard you purchase fits the case you have. For example, if you have an ATX case, you now know that an ATX or MicroATX motherboard can fit in that case.



Lab 1-3 will help you summarize distinguishing features of popular motherboard form factors. Lab 1-3 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Understanding Bus Architectures

The motherboard has a number of expansion slots that can expand the computer's capabilities. When the system is first purchased, a computer has only so many capabilities — the nice thing is that you can expand on those capabilities by purchasing cards to add to the expansion slots, or *bus architecture*.

Expansion slots expand on what the computer can do. The problem is that there are different types of expansion slots in the system, so when you go to purchase that sound card or network card, you have to make sure that you purchase the right type. In the following sections, I show you the different types of expansion slots and compare their characteristics.



Another term for the expansion slots is *bus architectures*. A number of different bus architectures have been developed over time. It is important to identify the differences between each of these architectures and also to know which ones are more popular today.

ISA

The *Industry Standard Architecture (ISA)* was the first major expansion bus architecture. It was originally developed as an 8-bit architecture and then evolved into a 16-bit architecture. The ISA bus architecture has a speed of 8 MHz, which is extremely slow by today's standards. Figure 1-36 shows two 16-bit ISA slots — note that the ISA slots are the black slots in the system.



Figure 1-36: Identifying ISA slots on the system.

ISA slots

One of the reasons why you still see 16-bit ISA slots in some earlier Pentium or Pentium II systems is because companies typically had a number of ISA network cards in the office from previous systems. When the company upgraded to the Pentium or Pentium II it was nice that they did not have to purchase new network cards because earlier Pentiums had ISA slots. Most systems today no longer have ISA slots. Figure 1-37 shows a 16-bit ISA network card.

MCA

One of the major downfalls of the ISA bus architecture is its performance. It runs at only 8 MHz, and it is only a 16-bit architecture — that was fine years ago, but everything evolves, and new and improved standards arise.




Figure 1-37:
Looking at a
16-bit ISA
network
card.

The *Micro Channel Architecture (MCA)*, which was developed by IBM, is a 32-bit architecture. The MCA architecture runs at 10 MHz and is not compatible with ISA. You usually find MCA slots in high-end IBM machines, such as those that might be used as a server.



Remember, ISA is an 8-bit or 16-bit technology that runs at 8 MHz. MCA transfers information in 32-bit chunks and runs at 10 MHz.

With MCA, IBM came up with a feature called *bus mastering*. Bus mastering works like this: Devices in the bus don't have to send information through the CPU if they want to talk to one another — they just send the information directly. This takes some of the workload off the processor and allows it to perform other tasks. Bus mastering became an important feature in future bus architectures. Figure 1-38 shows an MCA card.

EISA

In 1988, the industry standard for expansion cards was still ISA, but bus architectures had already been created that performed better. So a number of companies got together with the goal of extending ISA while maintaining backward compatibility so that companies could use their existing ISA cards.





Figure 1-38: Looking at an MCA network card.

As a result, the *Extended Industry Standard Architecture (EISA)* was developed as a 16- and 32-bit architecture. The big advantage to EISA is that it maintains support for the ISA cards that some companies already have in large quantities, and it also supports 32-bit EISA cards. EISA also included the major advancement in expansion bus technology that MCA created, known as bus mastering. Because both ISA and EISA cards fit into the same slot, they keep the same speed of 8 MHz.

The bus architecture holds both 16- and 32-bit cards because the EISA slots have two levels. The EISA cards have very deep edge connectors that fill the two levels (32-bit) of the slot, but ISA cards only fill the top level (16-bit). Figure 1-39 shows an EISA slot and the two different levels in the slot: one level for the ISA card to fill and the other level for the EISA card to fill.

Figure 1-39: Looking at how the EISA slot is organized.



VESA

In 1992, the *Video Electronics Standard Association (VESA)* developed a bus architecture that outperformed ISA. VESA is a 32-bit architecture that supports bus mastering and runs at the same speed as the processor, which, when VESA was created, was around 25 to 33 MHz. Because the bus runs at the speed of the processor, developers called this *VESA local bus*, or *VLB*. VESA slots are typically used for video cards.



Remember, EISA is an extension on ISA and is a 16-bit or 32-bit technology. For backward compatibility, EISA runs at 8 MHz. VESA is a 32-bit architecture that runs at the processor's speed. It is generally used for video adapters.

VESA slots are extremely easy to identify because they are tan and act as an extension to the ISA slot. You will notice the black ISA slots and then right beside them may be a tan slot. The VESA card fills the entire ISA slot and the additional extension to make the full 32-bit path for VESA. This allows an ISA card to be inserted into the slot for backward compatibility or, with the extension slot, the VESA slot can hold a VESA card. Figure 1-40 shows a VESA slot.



Figure 1-40: Looking at a VESA slot, which is an extension of the ISA slot.

PCI

Peripheral Component Interconnect (PCI) is one of the newer bus architectures to hit the market. PCI has two flavors: 32-bit cards and 64-bit cards. When Pentium systems hit the market, their motherboards featured both ISA/EISA slots and PCI slots. If you want to buy a new card today, you would most likely buy a PCI device for one of the PCI slots in your system.

The 32-bit version of PCI has a speed of 33 MHz, while the 64-bit version of PCI runs at 66 MHz. PCI also supports bus mastering. One of the other major benefits of PCI is that it is a Plug and Play architecture. If you are running a Plug and Play operating system like Windows 2000 or XP and your computer has a Plug and Play BIOS, then the system resources like IRQs and I/O addresses are dynamically assigned for PCI components.

PCI slots are easily identified on the motherboard as the small white slots, usually alongside the AGP slot. Figure 1-41 identifies PCI slots on a motherboard.




Figure 1-41: Installing a card into a PCI slot located on the motherboard.

PCI slots

PCMCIA

Personal Computer Memory Card Industry Association (PCMCIA) is a unique type of expansion bus architecture because of its small size. PCMCIA is popular in laptop computers. How are you going to get a big network card like the one that is used in a desktop computer into a little laptop to add network support? The answer is that you can't; you have to purchase a PCMCIA network card for the laptop to add network support. *PCMCIA cards*, also known as *PC Cards*, are a little bit larger than a credit card and can fit into your back pocket (though I don't suggest that you put one there). Figure 1-42 shows a PCMCIA network card.



Figure 1-42:
Looking at
a PCMCIA
network
card.

PCMCIA (say that five times fast!) is a 16-bit architecture that runs at 33 MHz and supports Plug and Play as well. Not only is PCMCIA a Plug and Play technology, but it is also a hot swappable technology. Hot swappable means that you can insert and remove PCMCIA cards without shutting down the system first. PCMCIA has three different types of slots named type 1, type 2, and type 3.

Table 1-1 shows the different PCMCIA slot types and the types of devices you can find in the different types of slots.

Table 1-1 Types of PCMCIA Slots

<i>Slot Name</i>	<i>Thickness</i>	<i>Types of Devices</i>
Type 1	3.3 mm	Memory cards
Type 2	5.0 mm	Modems/Network cards
Type 3	10.5 mm	Removable drives

Type 1 cards were originally used to add memory to laptop computers or personal computers. This is where the “personal computer memory card” part of the PCMCIA name comes from.



Remember, PCI is a 32- or 64-bit technology, runs at 33 MHz, and supports Plug and Play. PCMCIA is the expansion bus architecture used by laptop computers and is a 16-bit architecture that runs at 33 MHz.

AGP

Advanced Graphics Port (AGP) has been around since the Pentium II processor appeared in 1997. It’s a 32-bit bus architecture that runs at 66 MHz — which is twice the speed of the PCI bus. Today’s motherboards have one AGP slot to hold an AGP video card. The performance gain from the AGP port not only comes from the increase in speed, but also from the fact that the AGP bus has a direct path to the processor so that information travels quickly from the processor to the AGP card. Figure 1-43 shows an AGP slot beside some PCI slots.

AGP can run in different modes, and the different modes dictate the speed of the bus. 1x mode runs at 66 MHz (266 MBps), 2x runs at 133 MHz (533 MBps), 4x runs at 266 MHz (1.07 GBps), and 8x runs at 533 MHz (2.2 GBps)!

PCI-X

A fairly new bus architecture is the PCI-X bus architecture. Because PCI-X uses the same connector style as PCI, it is 100-percent compatible with PCI in the sense that it can hold PCI cards. So, a motherboard that has PCI-X slots can also house older PCI cards — that is a great feature!

Like PCI, PCI-X is a 32-bit and 64-bit bus architecture and is available in four different speeds. PCI-X runs at speeds of 66 MHz, 133 MHz, 266 MHz, and 533 MHz.




Figure 1-43:
Looking at
an AGP
card in an
AGP slot.

AGP slot

PCI Express

While PCI-X is compatible with PCI by being able to hold PCI cards and also sending data in parallel (multiple bits at one time), the PCI Express bus architecture takes a totally different approach. PCI Express is a serial bus that does not support existing PCI cards. The PCI Express slot, shown in Figure 1-44, is the smaller black slot and is much smaller than a normal PCI slot, so it can't possibly house a PCI card.

PCI Express slot




Figure 1-44:
Identifying a
PCI Express
slot on the
mother-
board.

PCI Express uses data lanes to transfer the information within the bus architecture. A data lane delivers an amazing transfer rate of 250 MBps per lane. PCI Express has different implementations, with each implementation having a different number of lanes identified by a multiplier. For example PCI Express with only one lane is known as x1 while a PCI Express bus with eight lanes is known as x8. The implementation of lanes allows PCI Express to reach fast transfer rates by implementing additional lanes. For example, current graphics cards for PCI Express have 16 lanes which provide a transfer rate of 4 GBps (16 x 250 MBps) — which is twice the rate of AGP 8x, which runs at 2 GBps).

There are currently systems with PCI Express at x1, x2, x4, x8, x16, and x32. The PCI Express slot gets bigger with each multiplier — for example, Figure 1-44 is displaying a PCI Express x1 slot which is the black slot only about one inch in length.

AMR and CNR

Audio/Modem Riser (AMR) is a newer bus architecture that adds a modem and audio card to the system. AMR allows the two components to be incorporated into a single card to reduce cost. Figure 1-45 shows an AMR slot on a motherboard.



Figure 1-45:
An AMR
slot on a
mother-
board.

AMR slot

Communication and Network Riser (CNR) is another new architecture that is used to implement LAN, audio, and modem functionality all in one.



As far as the “real world” and the exam are concerned, you need to be extremely strong in the area of bus architectures. A big part of servicing computers is installing network cards, sound cards, and video cards — these components come as ISA, PCI, or AGP cards today. You need to know how to look at a system and say, “We are going to buy a PCI network card for this system.”



Lab 1-4 will help you identify the different performance characteristics of each of the standard bus architectures. Lab 1-4 can be found in the `Labs .pdf` file in the Author directory of the CD-ROM.

Performance Considerations

When you want to improve a motherboard’s performance, one of the first things you should do is check the speed of the motherboard. For example, some systems today have 400 MHz motherboards, and some have 533 MHz motherboards. To find out the speed of the motherboard check the documentation for the board. You can get a performance increase from a faster motherboard.

Another performance consideration occurs when you add expansion cards to the system. You should first evaluate what expansion slots are free and then purchase a card that will give you the best performance. For example, if you need to buy a network card for your computer, you start by looking at what expansion slots are free. If only two slots are free, an ISA slot and a PCI slot, then your choices are limited to an ISA network card or a PCI network card. Because PCI outperforms ISA, however, you would be better off purchasing a PCI network card.

You can also get a performance increase from motherboards that have more cache memory. Look at the motherboard to see if there is a place to install some level 2 (L2) cache. L2 cache can dramatically increase performance because it is generally closer to the processor than the system memory (RAM) and is a faster type of memory than the system memory. Bottom line: the more cache memory you install, the better the motherboard’s performance.

Getting an A+

This chapter introduces you to a number of key components of the motherboard and different motherboard form factors. The following is a list of the key points to remember when dealing with motherboards:

- ◆ The *motherboard* (or *system board*) is the computer component that interconnects all other components.
- ◆ Serial (COM) ports come in two flavors: DB9-male and DB25-male. Parallel ports come in only a DB25-female port.
- ◆ There are two main types of cache memory: L1 and L2 cache. L1 cache memory is integrated into the processor, while L2 cache is contained outside the processor but in the processor casing or on the motherboard.
- ◆ IDE supports two devices in the IDE chain, whereas EIDE has two channels with two devices in each channel (a total of four devices).
- ◆ There are a number of major motherboard form factors: Full AT, Baby AT, NLX, and ATX, to name a few. Motherboard form factors differ in the size of the board and the layout of the components stored on the board.
- ◆ You may add components, such as a sound card or network card, to the computer by inserting an expansion card into one of the expansion slots in the system.
- ◆ ISA was the popular bus architecture for years, but because of its limitations (16-bit architecture and a speed of 8 MHz), it has been replaced by the PCI bus architecture. PCI is a 32-bit/64-bit architecture with a speed of 33 MHz.
- ◆ AGP is the common bus architecture used to insert a video card in today's systems.
- ◆ You may increase the performance of the system by using a faster motherboard or by purchasing better-performing expansion cards.

Prep Test

- 1** What was the original motherboard speed of the ATX board?
 - A 33 MHz
 - B 60 MHz
 - C 66 MHz
 - D 100 MHz

- 2** Which bus architecture supports 32-bit/64-bit cards and transfers information at 33 MHz?
 - A ISA
 - B EISA
 - C AGP
 - D PCI

- 3** Which motherboard component is responsible for charging the CMOS RAM so that CMOS can maintain its data?
 - A Battery
 - B BIOS chip
 - C CMOS chip
 - D Power supply

- 4** How many pins does a standard IDE controller have?
 - A 33 pins
 - B 40 pins
 - C 50 pins
 - D 20 pins

- 5** Which of the following best describes a Baby AT motherboard?
 - A Uses Slot 1
 - B Runs at 100 MHz
 - C The only I/O port is a keyboard port on the back
 - D Incorporates AGP

- 6** How many devices are supported in a USB chain?
 - A 10
 - B 27
 - C 127
 - D 255

- 7** How many pins does a standard floppy drive controller have?
- A 34 pins
 - B 40 pins
 - C 50 pins
 - D 75 pins
- 8** Which bus architecture might be found in older IBM servers?
- A ISA
 - B MCA
 - C VESA
 - D EISA
- 9** What type of cache memory will you find on a motherboard?
- A L1
 - B L2
 - C SDRAM
 - D SRAM
- 10** Which type of memory module supports 32-bit data chunks?
- A DIMM
 - B Cache
 - C SIMM
 - D Video
- 11** Which bus architecture runs at 8 MHz and supports 16-bit ISA cards?
- A MCA
 - B AGP
 - C PCI
 - D EISA
- 12** What is the bus architecture used in laptop computers?
- A PCI
 - B PCMCIA
 - C EISA
 - D ISA

Knowing Your Motherboard

- 13 Which of the following best describes AGP?**
- A AGP slots have a direct path to the processor to help increase performance of AGP devices.
 - B AGP cards run at 33 MHz.
 - C AGP runs at 66 MHz and gets access to the processor through the PCI bus.
 - D AGP stands for Advanced Graphics Port and is used to install additional video memory.
- 14 Which PCMCIA card type is used for modems?**
- A Type I
 - B Type II
 - C Type III
 - D Type IV
- 15 Which port is typically used for a modem?**
- A USB
 - B LPT1
 - C COM1
 - D LPT2
- 16 When connecting a floppy drive to the system, which end of the ribbon cable connects to the floppy drive?**
- A The end with a red stripe
 - B The end with a twist
 - C The end with a blue stripe
 - D The end without a twist
- 17 What are the labels given to the power connectors that supply power to the motherboard?**
- A P1 and P2
 - B P22 and P2
 - C PT1 and PT2
 - D PS1 and PS2
- 18 How many pins are in the end of a parallel cable connector that connects to the computer?**
- A 36
 - B 40
 - C 25
 - D 33

19 Which port sends information 8 bits at a time side-by-side?

- A COM1
- B LPT1
- C COM2
- D USB

Answers

- 1** D. The ATX board had an original motherboard speed of 100 MHz. Older boards, such as the Baby AT, had motherboard speeds of 60 and 66 MHz. *See "ATX."*
- 2** D. The PCI bus architecture is a 32-bit and 64-bit architecture that runs at 33 MHz. AGP runs at 66 MHz, while both ISA and EISA run at 8 MHz. *Review "PCI."*
- 3** A. The battery is responsible for maintaining a charge so that the CMOS RAM doesn't lose its information. The BIOS chip stores the core system code that allows all of the devices to communicate. *Check out "Battery."*
- 4** B. An IDE controller has 40 pins to allow a 40-wire ribbon cable to connect a hard disk or CD-ROM to the motherboard. A floppy controller uses 33 pins, and 50 pins are used by internal SCSI devices. *Peruse "IDE connections."*
- 5** C. The Baby AT motherboard uses a DIN connector as the keyboard connector and is the only I/O port found on the Baby AT and the Full AT motherboards. Both slot 1, which is the processor slot for Pentium II processors, and the AGP slot that is used by video cards, exist on the ATX board. *Take a look at "Baby AT."*
- 6** C. A computer can support up to 127 USB devices in a USB chain. *Peek at "USB ports."*
- 7** A. A floppy drive connector has 34 pins, while an IDE connector has 40 pins, and internal SCSI devices have 50 pins. *Look over "Floppy disk connectors."*
- 8** B. In the past, you were likely to see MCA in older IBM servers because IBM developed the MCA bus architecture. *Study "MCA."*
- 9** B. Level 2 (L2) cache is the type of cache memory that is found on motherboards, whereas Level 1 (L1) cache is found in the processor. Choices C and D are not types of cache memory. *Refer to "Cache memory."*
- 10** C. 72-pin SIMM modules are 32-bit modules, whereas a DIMM is a 64-bit memory module. *Examine "SIMM/DIMM sockets."*
- 11** D. The EISA bus architecture runs at 8 MHz and supports ISA cards. MCA does not support ISA, AGP runs at 66 MHz, and PCI runs at 33 MHz. *See "EISA."*
- 12** B. The bus architecture used in laptop computers is called PCMCIA (Personal Computer Memory Card Industry Association). The other three choices are bus architectures available to desktop computers. *Review "PCMCIA."*
- 13** A. The AGP slot runs at 66 MHz and has a direct path between the slot and the processor so that information will not have to travel through one of the slower buses. AGP stands for Advanced Graphics Port and is used to install a video card, not video memory. *Check out "AGP"*

- 14** B. Type II cards are used for network cards and modems. Type I cards are used for memory upgrades, and Type III cards are used for removable drives. *Peruse "PCMCIA."*
- 15** C. COM1 is a serial port that is typically used for modems or serial mice. LPT ports are parallel ports that are generally used to connect printers. *Take a look at "Serial ports."*
- 16** B. The floppy ribbon cable has one end that is twisted; that twisted end must be connected to the floppy drive. *Look over "Floppy disk connectors."*
- 17** A. P1 and P2 are the typical labels given to a motherboard's power connectors. *Study "Power connectors."*
- 18** C. A parallel cable has a different style of connector at each end of the cable. The end that connects to the computer has 25 pins, and the end that connects to the printer has 36 pins. *Refer to "Parallel port."*
- 19** B. Parallel ports send information in 8-bit chunks, side-by-side; LPT1 is the only parallel port listed. The other three ports are serial ports that send information 1 bit at a time. *Examine "Parallel port."*

Chapter 2: Picking Your Processor

Exam Objectives

- ✓ Understanding CPU characteristics
- ✓ Identifying popular CPUs
- ✓ Identifying sockets
- ✓ Installing a processor

Although all components of the computer function together as a team, every team needs a leader — someone who gives out instructions and keeps everyone working toward the same goal. If any PC component were to be considered the team leader, it would probably be the *central processing unit (CPU)*, also known as the *processor*. The key word here is *central*, which implies “center” or “focus.” The CPU can be considered the focus of the computer because it controls a large number of the computer system’s capabilities, such as the type of software that can run, the amount of total memory that the computer can recognize, and the speed at which the system will run.

In this chapter, you get a look at some of the features of the CPU that are responsible for regulating the capabilities of the computer system. It will also discuss the importance of the CPU and its role as a PC component, as well as identify some of the main characteristics that make one CPU better than another.



When preparing for the A+ exams, it is important that you’re comfortable with terms like *MMX*, *throttling*, and *cache memory*. You also want to be sure that you’re comfortable with the differences between various processors. For example, what makes a Pentium 4 better than a Celeron? Which AMD processor competes with the Celeron? All of these are questions you should know the answer to before you take the A+ Certification exams, and this chapter helps you find out the answers. Good luck!

Understanding Processor Terminology

In this section, you learn some basic terms that describe characteristics of different processors — past and present. The exam might not ask for the specific definition of each term, but understanding the terms will help you answer the related questions in this topic area.

Processor speed

Processor speed is the speed at which the processor executes its instructions or commands. This speed was originally measured in millions of hertz, or *megahertz (MHz)*, per second. A *hertz* is also known as a *clock cycle*, and a processor can execute code at every clock cycle. Thus, a processor operating at a measly 1 MHz per second can execute *one million tasks every second*. Processors today now measure their speed in *gigahertz (GHz)* per second. A *gigahertz* is a billion clock cycles per second — so the CPU can execute tasks a billion times per second!

Original CPUs had a speed of 4.77 MHz, while systems at the time of this writing are running around 3.0 GHz. Although processor speed is not the only factor affecting performance, in general, the faster the processor, the faster the system.

Data bus

A city bus is responsible for transferring people from one location to another. In the world of computers, a *bus* is responsible for delivering data from one location on the PC to another. The *data bus* is the term used to define the pathway between the processor and memory. Because the processor accesses information from memory so often, an entire bus — the data bus — is dedicated to this action. The larger the data bus, the more data can be carried from the CPU to memory in one clock cycle.

Here's an illustration. What would happen if 50 people needed to go from one end of the city to the other, but a city bus had only 25 available seats? The answer is simple. The bus would make two trips. But wouldn't it be more efficient to get a larger bus? If you upgraded the bus to 50 seats, the bus would have to make only one trip to transfer the 50 people from one end of the city to the other, which increases the efficiency of the public transit system.

The data bus works the same way, only it transfers data in the form of *bits* (a single bit is either a one or a zero). All data processed by the computer is in the form of bits. The data bus has a "full capacity" point at which it cannot handle any more bits of data, just as the bus system in the city has a "full capacity" point (measured in "seats").

If a processor has a 16-bit data bus, it means that it can deliver at most 16 bits during a single clock cycle. If the same processor needs to deliver 32 bits of information, it will have to take two trips, send 16 bits during the first clock cycle and the remaining 16 bits during the next clock cycle. Taking that same 32 bits of information and processing it on a 32-bit processor means that the information will be delivered in one trip — one clock cycle — as opposed to two, which increases the overall efficiency of the system.

Address bus

Figure 2-1 shows how system memory is organized like a spreadsheet, in rows and columns. These rows and columns make up blocks that can be written to and read from. If you want to store information in one of the blocks, you have to reference the location by address. For example, you may store data in cell B2.

	A	B	C	D
1				
2		↓		
3				
4				
5				

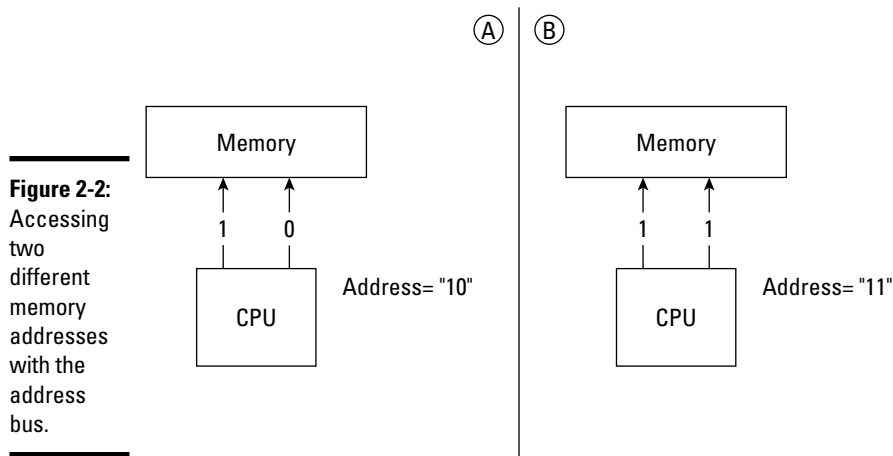
Figure 2-1:
How system
memory is
organized.

To store information into system memory, your processor has to give an address that points to a particular storage location, only the address doesn't look like "B2." It looks something like "10," or maybe "11," which are two completely different memory locations and, as a result, the data would get stored in two different blocks.

Your processor accesses memory locations through the *address bus*. If, for example, the address bus is two-bit, the processor has two address lines from the processor to system memory. The *address lines* carry signals that specify locations in memory, each with an on/off state. A "1" represents an on state, and "0" represents an off state. The combination of the on/off states of both address lines at any given time is how a reference to an area in memory is made. The left side of Figure 2-2 illustrates a processor making a reference, or *call*, to Address 10, while the right side shows a reference to Address 11. These two address calls reference completely different locations in memory.

If you add another address line to the address bus, the processor can access even more possible addresses because the processor has more variations with three bits than with two. A two-bit address bus can make a reference to four possible memory addresses (2×2), while a three-bit address bus can make a reference to eight possible memory addresses ($2 \times 2 \times 2$).

Therefore, the address bus dictates how much physical memory the processor can access. For example, an old 80286 processor has a 24-bit address bus, which means that it can access 16,777,216 (2^{24}) memory addresses, or 16MB of system memory. Newer processors have 36-bit address buses, which allows them to access 68,719,476,736 memory addresses, or 64GB of memory.



Registers

Registers are storage areas within the processor used to store data temporarily for manipulation later. They are used to store and process data and perhaps write back the result of the processed data. The benefit of storing this information in the registers instead of in memory is that the processor contains the information and does not have to retrieve it from memory — which takes time. It is as if information to be processed were in your pocket, rather than across a room, where you would have to walk all the way over and pick it up. Having information in your pocket means it can be accessed much more quickly, saving time and increasing performance. Registers give a processor quicker access to data, and the more registers a processor has, the more data it can store.

Registers are measured in bits. A processor with 16-bit registers has 16 containers into which a programmer can choose to store information, while a processor with 32-bit registers has twice as many containers that it can use to store information.

Cache memory

The processor accesses information that resides in system memory, which is a slower process than if the information is stored in the processor's own special "high-speed memory," known as *cache memory*. When the information is sitting in system memory and the processor sends a request for that information, the request goes to the *memory controller*, which manages data in memory. The memory controller finds the data in memory, retrieves it, and delivers it to the processor. Throughout this entire process, the processor is simply "waiting around" for the information. Thus, many of the newer processors include their own special high-speed memory within the processor's chip.

When the processor retrieves information from slower system memory, it then stores it in the high-speed cache in case the processor wants to access the information a second time. The benefit is that the second time the data is needed, it is sitting in the high-speed memory located on the processor chip. The processor will not need to sit around and wait for the data to come from system memory — again increasing overall performance.

Cache memory is integrated right into the processor's chip and is made up of *static RAM (SRAM)*. For more information on SRAM check out Book II, Chapter 3. Cache memory is very expensive because it is much quicker than regular system memory. As a result of this extra memory being integrated into the processor chip, the processor becomes more expensive than a processor that has less or no cache memory.

There are two types of cache memory: *Level 1 (L1) cache* and *Level 2 (L2) cache*. L1 cache is built into the processor, whereas L2 cache resides outside the processor. In the past, L2 cache resided on the motherboard, but newer processors have a bit of L1 and L2 cache in the chip package. If you upgrade the cache memory on your computer, you are adding L2 cache to the motherboard — you wouldn't be able to upgrade the L1 cache on the processor. Because L1 cache is built into the chip, you can't upgrade it without replacing the entire processor.

The integration of cache memory into processor chips didn't come to market until the 80486 chips were developed in 1989. Generally, 80486 chips had 8K of L1 cache, and the Pentium chip increased that amount to 16K. In fact, many of the newer processors have increased the L1 cache to over 16K and have also included some L2 cache. The more cache memory a processor has, the quicker (and more expensive) the system will be.

Math co-processor

The *math co-processor*, also known as the *Numeric Processing Unit (NPU)*, is the processor's sidekick. Systems that have math co-processors can well outperform systems that do not have math co-processors because the math co-processor takes some of the workload off the CPU. For example, it performs many of the large calculations that applications may require, such as floating point arithmetic. Overall system performance increases because the CPU can focus on logic functions while the math co-processor executes complicated mathematical functions.

If you have large spreadsheets or use large graphics applications, you may find that applications run very poorly or not at all on systems without a math co-processor. If you are running a system that does not have a math co-processor integrated into the CPU, then you can add one to the motherboard — or perhaps upgrade the main processor.

In earlier computers, the processor was one chip and the math co-processor was a separate chip on the motherboard. For example, years ago, a 386 computer used an 80386 chip on the motherboard as the processor, but you could add an 80387 chip to the board to act as the math co-processor. All processors since the 80486 computer, including Pentium-class systems, have a math co-processor integrated into the processor's chip, so you will not be adding a math coprocessor to the system.

Real-mode versus protected-mode

A *real-mode* processor is a processor that sees memory as a whole unit and deals with it as a single entity. In other words, if you have 512MB of RAM, the real-mode processor sees that as one block of memory. This is limiting because in order to run multiple programs at the same time, each program has to be assigned its own independent block of that 512MB — something that real-mode processors cannot do. As a result, real-mode processors don't have any *multitasking* capabilities — the capabilities to divide memory up into multiple parts and run different applications or tasks in each part.

Protected-mode processors support the segregation of system memory into different parts and assigning a different application to each part of memory. Therefore, protected-mode processors support multitasking and multitasking operating systems, such as Windows 2000 and Windows XP.

Protected-mode processors also support *virtual memory*, which is the process of using hard disk space as emulated memory. This means you could increase your 512MB of RAM by using 768MB of hard disk space as “pretend” RAM. In this case, as far as the applications that are running are concerned, the system has 1280MB of memory — the combination of true memory plus virtual memory.

MMX

After the Pentium was developed, Intel introduced a feature called *MultiMedia eXtensions*, or *MMX*. MMX added 57 new instructions that were built into the processor and told the system how to work with audio, video, and graphics. If these instructions were not built into the processor, the processor would have to retrieve them from somewhere else.

At the time MMX was developed, both the home and business user seemed to be heading toward the world of multimedia, and it made sense to enhance the processor and make it “multimedia-aware.” Running any kind of multimedia application on a processor that supports MMX gives you a major performance increase over a processor that doesn't support MMX technology.

Hyperthreading

Hyperthreading is a feature designed by Intel that was placed in the Pentium processors. *Hyperthreading technology*, or *HTT*, allows a processor to logically act as two different processors by being able to execute simultaneous threads. A *thread* is a part of an application that executes at any given time. For example, when running Microsoft Word, one thread accepts keystrokes, and another thread runs the spell checker while you type — two parts of the application run at the same time.

In order for a system to truly be able to take advantage of multithreaded applications, you normally need a system that has multiple processors — one processor to run one thread at a time. With hyperthreading, one processor is able to run more than one thread at a time, increasing performance by 15 to 30 percent.

Dual core processors

A *dual core* processor combines two independent processors and the L1 cache from those processors onto a single processor chip. The benefit of a dual core processor is that it can execute multiple threads at the same time without hyperthreading because you essentially have two processors.

A dual core processor has the benefit of having two processors' core features packaged into one physical processor. The core features include pipelines and cache memory. A dual core processor can benefit from the two processors' combined L1 cache on the same chip, meaning that each of the processors in the dual core each have a block of L1 cache available. The dual core processor also has a block of shared L2 cache between the two processors in the dual core chip.

A huge benefit of being only one chip on the motherboard is that the one dual core chip draws less power than two separate processors would. Figure 2-3 shows the logical view of a dual core processor.

Throttling

Throttling is a feature built into a lot of newer processors today and involves the CPU sensing when it is going to overheat and then reduces its speed to lower the heat to an acceptable range.

Processors that support throttling have a built-in thermal sensor (a high-tech thermometer) that monitors the temperature of the processor. When the processor detects that it is going to overheat, maybe due to a fan failure, the processor drops its speed so the temperature drops to an acceptable range.

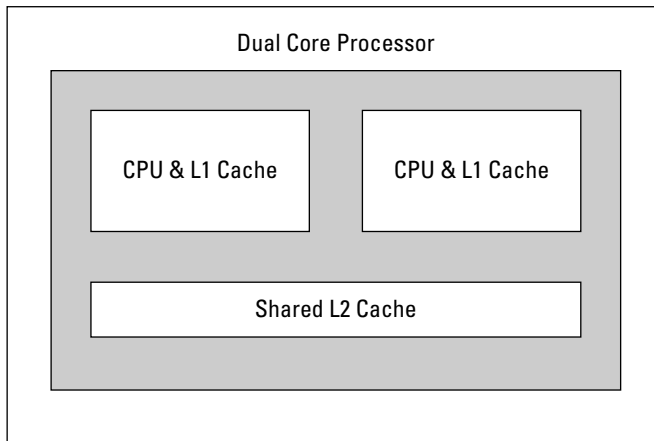


Figure 2-3: Looking at the logical structure of a dual core processor.

Overclocking

Overclocking is a big feature for PC enthusiasts and involves running a piece of hardware faster than the speed at which it is rated. A number of devices can be overclocked, such as video adapters and, of course, processors.



Although you may be able to overclock the processor, it is not recommended because overclocking can result in an unstable system or even hardware failure.

VRM

The *Voltage Regulator Module (VRM)* is responsible for regulating the voltage that is delivered to the processor. The VRM is located on the motherboard or appears as its own device in the system and provides the correct running voltage to the processor.

Some VRMs use a jumper on the motherboard to determine how much voltage is supplied to the processor, while other VRMs sense what the processor needs on startup. Typically, VRMs on the motherboard sense what voltage the processor needs and then supply that voltage.

Chip packaging

The term *chip packaging* refers to how the chip is constructed and delivered to the consumer. The chip package defines the appearance or form factor of the chip. Many chip packages have been used over the years.


The chip packages you should be familiar with for the A+ exam are as follows:

- ◆ **Dual Inline Package (DIP) chip:** A rectangular chip with two rows of 20 pins. Pin 1 is located at the end of the chip that has a square notch carved into it. It is important to identify Pin 1 because when you add a DIP chip to the motherboard, you will have to match Pin 1 on the chip with Pin 1 in the chip socket.

Older processors, such as the 8088 and many math co-processor chips, use the DIP chip style. Although they are no longer used for CPUs, DIP chips are still used for cache memory and BIOS chips on motherboards. They are also found on memory modules. (See Book II, Chapter 3, for a discussion of memory modules.)

- ◆ **Pin Grid Array (PGA) chip:** One of the most popular processor chip packages in use today, the PGA chip is a square chip that has an array of pins filling the shape of the chip. In general, the PGA chip uses hundreds of pins. You can locate Pin 1 on the PGA by identifying the corner of the PGA chip that has the corner cut off — that corner is where Pin 1 is located. Figure 2-4 compares a DIP (right side) with a PGA (left side) chip type.

Figure 2-4: Comparing a DIP chip package (right) to a PGA package (left).



Today's implementation of the PGA chip fits into a *Zero Insertion Force (ZIF)* socket. The ZIF socket is ideal for upgrading processors compared to the days before ZIF sockets were used because the ZIF socket has a lever on the side of the socket that you pull up on, which raises the chip out of the socket. Because the chip is automatically raised out of the socket, it allows you to simply remove the chip out of the socket with little effort! Before ZIF sockets were used, you had to pry the chip out of the socket trying to ensure that you did not damage the chip or the pins. With the ZIF socket, after the processor is raised, you can replace the old chip with a new one. In the past, not all boards used ZIF sockets, so you had to get some special extractors to pull the chip out (carefully!). Figure 2-5 shows a ZIF socket.




Figure 2-5:
A ZIF socket
on the
mother-
board holds
a processor.

- ◆ **Single Edge Contact (SEC) chip:** A chip package type that was popular with the Pentium II processors, the SEC chip is a huge cartridge surrounded by a plastic casing. The newer version, SEC2, is implemented as a card that is inserted into a slot on the motherboard and doesn't have the big plastic casing around it. It is important to stress that the SEC and SEC2 are inserted into a slot and not a socket. For more information on slots and sockets, read the next section. Figure 2-6 shows an SEC chip package along with some PGA chip packages.



Be sure to remember the different chip package types for the A+ exams. The Pentium II processor used the SEC, while the newer processors such as the Pentium 4 are using the PGA.




Figure 2-6: An SEC chip package along with some PGA chip packages.

Identifying Socket Types

Intel decided to develop a new standard for upgrading a processor on motherboards, beginning with the 80486 chips and continuing with the Pentium-class processors. This standard was called processor sockets. A *processor socket* is a socket designed to hold a specific processor chip with the appropriate number of pins. This enabled Intel to develop new chips with compatibility of a particular pins. For example, if a socket is developed with 321 pins, Intel could develop a new processor that has 321 pins and know that the processor will work with any motherboard that has the right socket. This allows the consumer to upgrade a processor much easier than in the past. Intel could design a new chip for an old socket so that customers could update their computers by dropping the new processor in the compatible socket.

Original Pentium processors supported mainly Socket 5 with 320 pins or Socket 7 with 321 pins. Thus, to add a Pentium processor to a motherboard, you would have to find out what socket existed on that board and then

purchase a CPU that would fit in that socket. You would also have to remember to match the voltage of the board to the voltage required by the CPU. Figure 2-7 will help you identify a CPU socket in your system.

The sockets are normally labeled with the type of socket it is along the side of the socket. For example, notice in the figure that the socket is labeled as PGA 370, meaning it's Socket 370 and will hold any processor designed for socket 370. Socket 370 is a socket that holds a processor containing 370 pins.





Figure 2-7:
Identifying a
processor
socket.

Table 2-1 lists the different types of sockets and the processors that are placed in the sockets. For more information about the processors, read the sections, “Looking at Popular Intel Processors” and “Don’t Forget Non-Intel Chips,” later in this chapter. Table 2-1 also shows the number of pins associated with the different types of sockets.

Table 2-1 Processor Socket Types		
<i>Socket</i>	<i>Processor</i>	<i>Number of Pins</i>
Socket A	Later Athlon, Duron, and Athlon XP	462
Socket 1	80486, 80486DX2, 80486DX4	169
Socket 2	80486, 80486DX2, 80486DX4	238
Socket 3	80486, 80486DX2, 80486DX4	237
Socket 4	Pentium 60/66	273
Socket 5	Pentium 75-133	320
Socket 7	Pentium 75-200	321

<i>Socket</i>	<i>Processor</i>	<i>Number of Pins</i>
Socket 8	Pentium Pro	387
Socket 370	Celeron and Pentium III	370
Socket 418	Itanium	418
Socket 423	Pentium 4	423
Socket 478	Later Celerons and Pentium 4	478
Socket 603	Xeon (Pentium 4 version)	603
Socket 611	Itanium	611
Socket 940	Opteron	940
Slot A	Athlon	242
Slot 1	Pentium II and Pentium III	242
Slot 2	Xeon	330



It is important to know the socket types used to hold the Pentium II, Pentium III, Pentium 4, Celeron, Athlon, Athlon XP, and Duron processors. You will not be expected to memorize the entire chart, but you should be familiar with the sockets used by today's popular processors.

Originally, the sockets were simply called Socket 1, Socket 2, and so on up to Socket 8. To make it easier to understand what processors went into which sockets, Intel started naming the sockets after the number of pins that existed on the processor that the socket would support. For example, Socket 370 holds a processor with 370 pins, while Socket 478 holds a processor with 478 pins. It is much easier now to identify what processors go into which sockets!

Now that you understand some of the characteristics of processors and you understand what a socket is, take a look at some of the popular Intel and AMD chips you are expected to know for the A+ exams.

Looking at Popular Intel Processors

In this section, I provide an overview of the Pentium-class processors and their characteristics, including data bus, address bus, registers, and the amount of cache memory supported on these processors. You will also be introduced to any new or unique processor features that each processor offers.

Pentium

The original Pentium processor was released in 1993 and was developed at speeds of 60 MHz and 66 MHz. The Pentium processor was a PGA chip that was placed in Socket 5 or Socket 7. Soon after its release, Intel marketed

Pentium processors in 75 MHz, 90 MHz, 100 MHz, 120 MHz, 133 MHz, 150 MHz, 166 MHz, and 200 MHz flavors, which were really just clock multipliers of the original 60 MHz or 66 MHz systems.

Clock multiplying is the concept that the processor will run faster than the motherboard that the processor sits in. For example, the original Pentium processor ran on 60 or 66 MHz motherboards. Say that the computer is marketed as being a Pentium 90. Since we know that the motherboard runs at 60 or 66 MHz, we can determine that the 90 comes from $60 * 1.5$ — meaning that the processor runs 1.5 times the speed of the motherboard. This is important because, as a consumer, when you purchase a computer, you want to make sure you know what the motherboard speed is, too, not just the advertised speed of the processor.

From a consumer's point of view, clock multipliers become important when you take a look at computers such as the Pentium 133 and the Pentium 150. Which is faster? The obvious answer is the Pentium 150, the system with the higher megahertz speed. But is it really? The Pentium 133 is a clock double of the 66 MHz board, while the Pentium 150 is a clock double and a half of the 60 MHz board. My point being that the overall performance of the system is controlled by more than just the speed of the processor — you need to consider other components such as the speed of the motherboard.

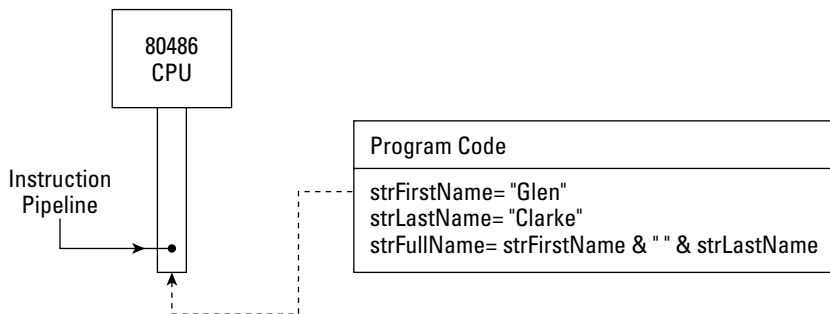
By looking at the motherboard speeds of the Pentium 133 and the Pentium 150, you could assume that the computer running the Pentium 133 may be able to keep up with, if not outperform, the one running the Pentium 150. Table 2-2 compares the speed of the motherboard and processor for the different Pentium systems.

<i>Processor</i>	<i>Motherboard Speed (MHz)</i>	<i>Multiplier</i>	<i>Processor Speed (MHz)</i>
Pentium 90	60	1.5	90
Pentium 100	66	1.5	99
Pentium 120	60	2	120
Pentium 133	66	2	132
Pentium 150	60	2.5	150
Pentium 180	60	3	180
Pentium 200	66	3	198
Pentium II	100	4.5	450

The Pentium processor has a 32-bit address bus, 32-bit registers, and a 64-bit data bus. It also has 16K of L1 cache that is divided into two 8K channels. One channel is for data cache and the other for application code cache.

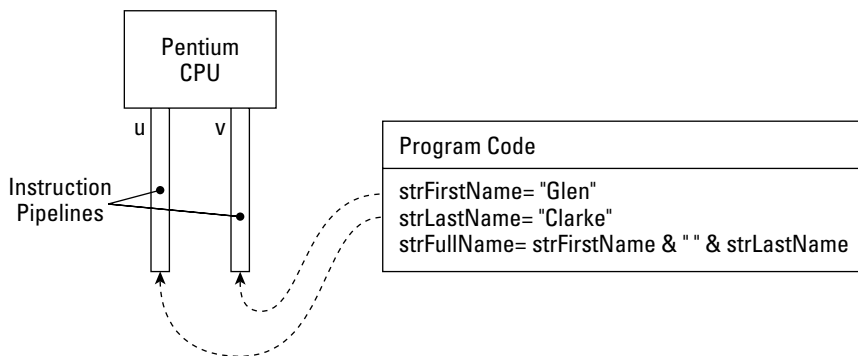
Before the Pentium came along, processors used one *instruction pipeline*. This meant that when an application executed, it would run each stage of the application job one step after the other. For example, if an application has three lines of code, as seen in Figure 2-8, each line of code can only be processed after the previous line of code is fully completed. This creates a delay, or wait time, that slows performance.

Figure 2-8: Single instruction pipelined processor executing application code.



The Pentium processor introduced a feature called *superscalar design*, which is the fact that the processor has two instruction pipelines, named U and V. Having two instruction pipelines enables the processor to execute two instructions at the same time. Thus, the three lines of program code, shown in Figure 2-9, can be quickly executed on a Pentium processor because Lines 1 and 2 are processed at the same time, causing Line 3 to be processed that much sooner. Notice that Lines 1 and 2 execute parallel to one another; therefore, *parallel processing* is taking place.

Figure 2-9: Dual instruction pipelined processor processing application code.



An application has to be designed to take advantage of two instruction pipelines. These applications are often labeled something like “Pentium Aware” or “Pentium Ready.”

Pentium Pro

In 1995, Intel released the Pentium Pro chip, which added a new level of performance to the Pentium processor. The Pentium Pro had all the characteristics of the Pentium processor — such as a 64-bit data bus and 32-bit registers — but it increased the address bus to 36 bits, which means that the Pentium Pro can access 64GB of RAM. The speed of the Pentium Pro ranges from 120 MHz to around 200 MHz.

The Pentium Pro includes two additional features on its chip that help it outperform the original Pentium. First, the Pentium Pro chip is really a two-chip team. One chip was the actual processor (with 16K of L1 cache, like the Pentium chip), but the other chip holds an extra 256K of cache memory. Since this cache memory is physically outside of the CPU, it is considered L2 cache.

The second feature that leads to the performance gain of the Pentium Pro is what is known as *dynamic execution*. Dynamic execution has three stages: multiple branch prediction, dataflow analysis, and speculative execution.

- ◆ **Multiple branch prediction** is the idea that the processor will look ahead and predict a number of instructions that may be needed in the very near future.
- ◆ **Dataflow analysis** occurs when the processor looks at the instructions it has predicted will be needed next and then assigns them a logical order of execution.
- ◆ **Speculative execution** is the actual execution of a given instruction based on the prediction and the order of execution assigned.

The Pentium Pro chip, shown in Figure 2-10, was implemented as a PGA chip that was placed in Socket 8.




Figure 2-10:
The Intel
Pentium Pro
processor.

Pentium II

In 1997, Intel produced the Pentium II, which was really just an enhanced Pentium Pro with speeds ranging from 233 MHz to 450 MHz. The Pentium II had a 64-bit data bus, a 36-bit address bus (64GB of RAM), and 64-bit registers and supports features such as MMX.

The Pentium II increased the amount of L1 cache that was integrated into the CPU to 32K, as opposed to 16K. The 32K of L1 cache was still divided into two equal channels: one 16K channel for data and one 16K channel for application code.

Intel packaged the Pentium II in the *Single Edge Contact (SEC)*, sometimes also referred to as the *Single Edge Contact Connector (SECC)*, that fits into Slot 1 on the motherboard. The SEC is a module enclosed in a casing or shell with two chips inside, one chip being the processor and the other chip being the 512K of L2 cache. Refer to Figure 2-9 to see what a Pentium II processor, which uses the SEC, looks like.

Another enhancement that accompanied the Pentium II was *Single Instruction Multiple Data (SIMD)*. To visualize how SIMD works, imagine five toddlers in a playroom, and that these toddlers are at the entertaining age of two — the age, of course, when the toddlers are preparing for their teen years by answering “no” to everything you say. You walk into the playroom and see that the five toddlers have found your box of darts and are throwing them at the walls. You are faced with a choice: You can either walk around to each child and explain why throwing darts at your walls is not a good idea (which means you will have to explain the same thing five different times), or you can have a good scream at the top of your lungs, which means that all the children will stop immediately and listen. SIMD works on the same basic principle. With SIMD, the processor gives the instruction to multiple processes at once — instead of having to give the same instruction multiple times. Thus, the processor saves time and creates a much more efficient way to work with information.


Celeron

The Pentium II processor performs very well, and with all that cache memory, it should! Unfortunately, that performance comes with a price. If you are not willing to pay that price, Intel has created a chip for you: the Celeron chip!

The Celeron chip is nothing more than a less-expensive version of the Pentium II processor with the built-in L2 cache either removed entirely or reduced. The first-generation Celeron chip was code-named the Covington; it has no L2 cache memory on it. The second-generation Celeron was code-named the Mendocino, and it contains 128K of L2 cache. Although this version of the Celeron does have L2 cache, it is dramatically reduced from the Pentium II's 512K so that it can be sold at a lower price.

The original Celeron shipped in an SEC package but also had a version that was packaged as a PGA, as shown in Figure 2-11.

Figure 2-11:
Intel's Celeron processor was first implemented as an SEC package, but later had a PGA chip that was placed in Socket 370.



Pentium III


The Pentium III processor shares many of the Pentium II's characteristics. It supports dynamic execution (as the Pentium Pro also did) and MMX technology, has 32K of L1 cache, and has either 256K or 512K of L2 cache. The Pentium III runs at a speed of 450 MHz to 1000 MHz, or 1 GHz.

The Pentium III chip offers 70 additional instructions that are integrated into the chip, enhancing the user's experience with 3-D graphic applications. The Pentium III chip also supports a number of low-power states to help conserve energy when the system is not in use. This processor is designed to run on either 100 MHz or 133 MHz motherboards.

Also note that there is a Pentium III version of the Celeron chip that runs as fast as the Pentium III processor but again has the L2 cache memory reduced. So now there are multiple versions of the Celeron chip — the PII version and the PIII version.

The Pentium III processor shipped in the SEC2 package (shown in Figure 2-12) originally, but was then packaged as a PGA chip. The SEC2 goes in Slot 1, while the PGA chip is inserted into Socket 370.

Figure 2-12:
The
Pentium III
processor in
the SEC2
package
that lives in
Slot 1.



Xeon

The Xeon processor is built on the Pentium II and Pentium III architecture — meaning that, like the Celeron, there is a PII version and PIII version of the Xeon. The Xeon chip is designed for higher-end systems, such as server-class systems, and contains more cache memory than the typical PII and PIII. The Xeon comes in flavors of 512K, 1MB, and 2MB of L2 cache.

The Xeon can also address 64GB of RAM and is designed for multiprocessing systems. A *multiprocessing system* is a computer with a motherboard that supports multiple CPUs. The Xeon processor has been designed to coexist with two, four, or eight CPUs.

The Pentium II Xeon and Pentium III Xeon chips were originally packaged as an SEC (shown in Figure 2-13) that was placed in Slot 2, but later versions use the PGA and are placed in Socket 603. The Xeon chip also contains a thermal sensor that shuts the processor down if it starts to overheat.

The Celeron is a scaled-down version of the Pentium II or III processor, and the Xeon is a step up from the Pentium II or III. There are also PIV (Pentium 4) versions of the processors:

- ◆ **PIV XEON:** Designed to work with a multiprocessing system that uses one or two processors.
- ◆ **PIV XEON MP:** Designed to work with a multiprocessing system that uses four or eight processors.

Figure 2-13: The Pentium II Xeon processor was originally delivered in the SEC package.



Pentium 4

The Pentium 4 processor runs at between 2 GHz and 4 GHz. The Pentium processor has 20K of L1 cache and 512K of L2 cache. The processor is shipped as a 423-pin or 478-pin PGA package, which means that the chip will be placed in Socket 423 or Socket 478 (shown in Figure 2-14).

The Pentium 4 processor gets a huge performance benefit by being able to perform four data transfers in one clock cycle along the *front side bus (FSB)*. The FSB is the bus that connects the processor to system memory (see Chapter 1 of this minibook).

Itanium and Itanium II

Intel created its first 64-bit processor in the Itanium and Itanium II processors. Because they were designed as 64-bit processors, you will be able to run 32-bit code on them, such as most copies of Windows and Office applications, but you will not be leveraging the 64-bit architecture by running 32-bit code. Special 64-bit editions of Windows can run on the Itanium processor, which enables you to take advantage of the 64-bit architecture. To learn more about the 64-bit editions of Windows, check out www.microsoft.com/windowsxp/64bit.

The original Itanium processor used a special packaging known as the *Pin Array Cartridge (PAC)*, which uses 418 pins, while the Itanium II was packaged in *Organic Land Grid Array (OLGA)* — which is a variation of the PGA, but the chip is located on a *processor card* (a circuit board that holds the processor). The OLGA fits into Socket 611.




Figure 2-14: Socket 478 can house a Pentium 4 processor.

The Itanium processor runs at around 1 GHz and contains a large block of cache memory: 32K of L1 cache, 96K of L2 cache, and 2MB or 4MB of L3 cache. The L3 cache is an additional block of cache memory located in the chip packaging.



Moving from 32-bit processors and applications to 64-bit versions would truly benefit any user that is using applications that are memory-intensive or calculation-intensive. For example, a user who works a lot with multimedia-type applications would see an improvement in performance.

Pentium “M”

For years, laptop manufacturers have been asking for smaller processors to place in laptop systems, and they finally have their wish. A number of processors have come out with the “M” version, which stands for *mobile*. The mobile version of the processors are smaller than the processors that go in desktop systems, so they will fit better and also use a lot less power. The benefit of using less power also means that they run much cooler.



Because the mobile versions of the processors use less power, they also are going to run a little slower than their desktop counterparts.

Some popular brands of mobile processors are the Intel Pentium III M and the Pentium M. Intel's big competitor, AMD, also has mobile versions of their processors: Athlon XP M and Mobile Duron. (Some manufacturers put the word *mobile* in the name of the processor instead of the letter *M*.) The next sections discuss more about AMD processors.

Don't Forget Non-Intel Chips

One of Intel's major competitors is *Advanced Micro Devices, Inc.* AMD has developed a family of processors that compete with the Pentium-class processors. In this section, I provide an overview of some of the characteristics of the AMD processors.

K6

The AMD K6 processor was designed to compete with the original Intel Pentium. The K6 has 64K of L1 cache, supports MMX technology, and has built-in branch prediction techniques. This processor has 321 pins, which means that it will fit into a Socket 7–supported motherboard.

K6-2

The K6-2 processor was designed to compete with the Pentium II chip. It has 64K of L1 cache and 256K of L2 cache. The K6-2 also supports dynamic execution, MMX technology, and superscalar design.

The K6-2 has added 3DNow! Technology — a number of additional instructions integrated into the chip to improve 3-D graphics applications. The K6-2 chip also uses a 100 MHz motherboard speed, which is a big improvement over the 60/66 MHz motherboard speed that the original Pentiums were using.

The K6-2 has 321 pins, which means that it will fit into a Socket 7–supported motherboard.

K6-III

The K6-III processor is designed to compete with the Pentium III chip. This chip shares many of the features of the K6-2, including a 100 MHz system bus. One of its new features is a Tri-Level cache. Not only can it take advantage of an L1 and L2 cache but also an L3 cache that can be included on the motherboard.

Athlon

The AMD Athlon chip has 128K of L1 cache and 512K of L2 cache. It supports improved dynamic execution, MMX technology, and 3DNow! Technology. The Athlon chip runs at speeds of up to 1.2 GHz and is designed to run on a 200 MHz system bus speed.

Unlike the K6-2 and K6-III, the Athlon is not a PGA-packaged chip that supports Socket 7. It uses its own socket type, called *Slot A*, because the processor is packaged as an SEC. The Slot A socket is not compatible with Intel's Slot 1, which means users have to purchase a motherboard designed for the Athlon chip.



Later versions of the Athlon moved to the PGA package that has 462 pins. These PGA chips are placed in Socket A.

Athlon XP

After the Athlon chip was produced, Intel created the Pentium 4 chip. So AMD wanted to create a competing chip for the Pentium 4, the Athlon XP. The Athlon XP is packaged as a PGA with 462 pins and is placed in Socket A. The Athlon XP runs at 2 GHz or more and contains 128K of L1 cache and 512K of L2 cache.



AMD markets these processors a little differently. Instead of labeling the processor with its speed, AMD labels it with its competitor's speed. For example, the Athlon XP 1800+ is rated at 1.6 GHz but runs as fast as Intel's 1.8 GHz processor.

Duron

AMD wanted to create a processor that competed with each version of the Intel processors. So, if the Athlon XP competes with the Pentium 4, what competes with the Celeron? You guessed it — the Duron.

The Duron has 128K of L1 cache and 64K of L2 cache. This processor is packaged as a PGA with 462 pins, which means it too goes into socket A.

Opteron

Just as the Duron was built to compete with Intel's Celeron, AMD created the Opteron to compete with Intel's 64-bit Itanium processors. The Opteron runs at about 1.8 GHz and contains 128K of L1 cache and 1MB of L2 cache.

The Opteron is packaged with a Micro-PGA, which is made up of 940-pins and is placed in Socket 940. One of the major differences between the Opteron and the Itanium is that the Itanium cannot run 32-bit applications; AMD decided that the Opteron would run in a 32-bit *or* 64-bit mode, thus allowing it to run 32-bit applications.

Installing a Processor

Now that you understand some of the popular processors that exist today, take a look at how to install a processor. This section identifies installation decisions you have to be aware of before actually attempting to install the processor.

Will it fit in the socket?

The first thing you need to verify before you purchase a new processor for your system is what socket type you have on your motherboard. You want to make sure that you purchase a processor that fits in that socket. For example, if you have Socket A on the motherboard, what processors fit in Socket A? If you said Athlon, Athlon XP, and Duron, you are correct.

Also be sure you know how many pins the socket has, because some processors support a few different-size sockets. For example, Intel makes both Socket 423 and Socket 478 versions of the Pentium 4, so you need to make sure you get the correct version of the Pentium 4 for your socket.

CPU voltage and transistor integration

Another important CPU characteristic that you have to watch for when upgrading your processor is the voltage the processor requires. The *voltage* is the power the processor draws from the main motherboard, which the motherboard receives originally from the power supply.

A processor is designed to run at a certain voltage. You need to ensure that the motherboard you are placing the processor into provides that voltage. If a motherboard supports more than one voltage, you can typically change a jumper on the motherboard, which will then control the voltage used by the processor. For more information on jumpers check out Book II, Chapter 1.

Performing the installation

Because most systems today are using ZIF sockets and PGA chips, I will discuss installing a processor into the ZIF socket. After you have verified that

your new processor will work with your motherboard, you are ready to install the processor. To install the processor, first remove the existing one by pulling up on the lever on the ZIF socket. When you pull the lever on the ZIF socket, the existing processor should rise out of the socket a bit.



Be sure to ground yourself before touching the insides of the computer. It is a great idea to get an antistatic wrist strap and clamp it to the computer's chassis so that you have a constant ground. For more information on safety procedures, refer to Book I, Chapter 3.

When the processor has risen a bit out of the socket, you can then gently lift the processor out (as shown in Figure 2-15). Be sure to lift the processor straight up so that you do not bend any of the pins.




Figure 2-15: Removing the processor from its socket.

After you have the old processor out of the socket, you can install the new processor by first finding out where Pin 1 is on the processor chip. Pin 1 is located in one of the corners of the chip and is usually indicated with a gold line marked on the bottom of the chip that contains the pins. If you don't see a line indicating where Pin 1 is, you will notice that one of corners of the square PGA is cut off (see Figure 2-16) — this corner is Pin 1.




Figure 2-16: The cut-off corner of the processor indicates the location of Pin 1.

After you have located Pin 1 on the PGA chip, you also need to figure out where Pin 1 goes in the socket. Again, you can figure this out by finding the “cut-off” corner of the socket. This corner is where the cut-off corner of the processor goes, as seen in Figure 2-17.

When you have matched up Pin 1 on the PGA chip with Pin 1 on the ZIF socket, carefully place the processor into the socket and then pull the lever down to lock it in place.



Just lay the chip into the socket; don't push it in. The whole point of a *zero insertion force* socket is that you don't have to risk damaging the pins by applying pressure.

Now that you have the processor in the processor socket, you need to install something to keep it cool, such as a heat sink or fan — or maybe even both.




Figure 2-17: Identifying the cut-out corner in the processor socket.

Pin 1 indicator

Keeping a Processor Cool

Processors are made up of thousands, even millions, of transistors. A *transistor* acts as a switch, either permitting or prohibiting the flow of electrical current. If current is allowed to flow through the transistor, some result is generated. If the current is not allowed to flow through the transistor, a different result is generated.

A processor contains millions of transistors that each hold an electrical charge, causing the processor to run at very high temperatures. Therefore, it

is important that you keep the processor cool. The most common cooling mechanisms today are heat sinks and CPU fans, which are sometimes used in tandem.

A number of other cooling devices are on the market today, and they are a little more expensive than your typical heat sink or CPU fan. The following are other cooling techniques you may find in systems today:

- ◆ **Liquid cooling:** A liquid cooling system pumps a cooling liquid throughout the PC by using small hoses. The benefit of a liquid cooling system is the reduced noise, but its big drawback is the amount of space needed in the PC for the components of the cooling system and, of course, the threat of a leak if the cooling system is not installed properly.
- ◆ **Temperature sensors:** A number of processors today come with a built-in thermal sensor (a high-tech thermometer). Temperature sensors allow the processor to identify that it is overheating and shut itself down until the temperature drops to normal.
- ◆ **Thermal compound:** This is a liquid paste that is placed between the processor and the heat sink to help draw the heat away from the processor and pass it through the heat sink.

Heat sinks and CPU fans

Due to the size of the Pentium processor and the number of transistors passing current, the chip can get so hot that it becomes unstable. Thus, many Pentium processors come with either a cooling fan or heat sinks. A number of processors today have a heat sink with a fan on top of the heat sink.

Heat sinks are a group of metal pins that are placed on the chip to draw heat away from it. A *cooling fan* is a small fan placed on top of the processor to pull the hot air away, helping to keep the processor cool. Figure 2-18 shows a heat sink.

Installing a heat sink and fan

Some processors may get so hot that a heat sink may not be enough of a cooling device; in this case, you may want to place a fan on top of the heat sink. To install the heat sink and fan on your system, simply place the heat sink on the processor and then clamp it in place with the heat sink clamping bar. After you have the heat sink in place, you can secure a fan on top of it by clamping the fan on the heat sink, as shown in Figure 2-19.

The term *passive heat sink* is used for a heat sink that does not use a fan on top, while the term *active heat sink* is used for a heat sink with a fan on top.






Figure 2-18:
Looking at a
heat sink.




Figure 2-19:
Placing the
fan on top of
the heat
sink.

Increasing Performance

When it comes to processors, there are a number of different ways to increase the performance of your system. A first and obvious way is to buy the faster processor when upgrading; for example, upgrade a 1.8 GHz processor to a 3

GHz processor if possible. Also, get a processor that is designed to run on the faster motherboards. For example, back when the Pentium II processors were popular, there were 100 MHz motherboards or 133 MHz motherboards — you get a faster system by having a 133 MHz motherboard.

You will have to look at other features of the processor, such as the L1 cache and L2 cache that resides in the processor packaging. Acquiring a processor with more cache memory can dramatically increase system performance.

Getting an A+

This chapter provides an overview of the key terms that are used to identify the popular processors and their capabilities. Some of the points you need to remember when preparing for the exam are:

- ◆ The three major chip packages are DIP, PGA, and SEC. PGA being the popular chip packaging used in today's systems.
- ◆ The speed of the processor is measured in Gigahertz (GHz), but has been measured in Megahertz (MHz) in the past.
- ◆ L1 cache is cache memory integrated into the processor chip, while L2 cache is found outside the CPU chip.
- ◆ A socket is used to hold the processor in place on the motherboard. Be sure to be familiar with the sockets for Intel's Pentium III, Pentium 4, and Celeron chips. Also know about the sockets for AMD's Athlon, Athlon XP, and Duron chips.
- ◆ Be sure to review the characteristics of different Intel chips and AMD chips.

Prep Test

- 1 Which of the following best describes superscalar design?**
 - A The processor is designed using only 3.1 transistors.
 - B The processor predicts the next few instructions to be executed and then determines the optimal order for the execution of these instructions.
 - C The processor has two instruction pipelines, which enables multiple instructions to execute at the same time.
 - D The processor works twice as fast as the motherboard.

- 2 What socket/slot would you find a Pentium II processor in?**
 - A Socket 5
 - B Socket 1
 - C Socket 7
 - D Slot 1

- 3 Which of the following are placed in Socket A? (Select all that apply.)**
 - A Pentium III
 - B Athlon XP
 - C Celeron
 - D Pentium 4
 - E Duron
 - F Itanium

- 4 What chip type was the original Pentium processor packaged in?**
 - A SEC
 - B PGA
 - C DIP
 - D Socket 5

- 5 Which of the following is a characteristic of protected-mode processors?**
 - A Heat sink support
 - B Virtual memory support
 - C Run only one application at a time
 - D Encased in protective shell

Picking Your Processor

- 6** How much L1 cache does a Pentium III processor have built in?
- A 32K
 - B 64K
 - C 8K
 - D 16K
- 7** Which of the following acts as a storage container for information that will be processed by the processor?
- A Data bus
 - B Address bus
 - C Registers
 - D Math co-processor
- 8** How much memory can a Pentium Pro address?
- A 128MB
 - B 512MB
 - C 4GB
 - D 64GB
- 9** Which statement best describes the purpose of a math co-processor?
- A The math co-processor performs all of the logic functions on behalf of the processor.
 - B The math co-processor performs floating point calculations on behalf of the processor.
 - C The math co-processor runs all applications, while a processor runs the operating system in a multitasking environment.
 - D The math co-processor allows for communication between devices.
- 10** Which sockets/slots do original Pentium chips typically fit into? (Choose two.)
- A Socket 1
 - B Socket 5
 - C Slot 1
 - D Socket 7
- 11** How much L1 cache memory does a Pentium II have built in?
- A 8K
 - B 16K
 - C 32K
 - D 64K

- 12** What sockets do Pentium 4 processors fit into? (Choose two.)
- A Socket 370
 - B Socket 423
 - C Socket 478
 - D Socket 920
- 13** What is the major difference between a Celeron processor and a Pentium 4 processor?
- A The Celeron has more L1 cache memory.
 - B The Celeron has less L1 cache memory.
 - C The Celeron has more L2 cache memory.
 - D The Celeron has less L2 cache memory.
- 14** Which of the following CPU characteristics determines how much total memory the system can access?
- A Data bus
 - B Address bus
 - C Registers
 - D Math co-processor
- 15** What chip package type was the Pentium II processor packaged in?
- A Slot 1
 - B PGA
 - C DIP
 - D SEC
- 16** How many instruction pipelines does a Pentium processor have?
- A 1
 - B 2
 - C 3
 - D 4
- 17** What chip type uses a ZIF socket?
- A SEC
 - B DIP
 - C PGA
 - D Socket 5

Picking Your Processor

18 Which processor runs on a 133 MHz motherboard?

- A Pentium
- B Pentium II
- C Pentium III
- D AMD K6

19 A Pentium 133 runs on what speed motherboard?

- A 60 MHz
- B 66 MHz
- C 100 MHz
- D 133 MHz

20 What type of cache is integrated into the Pentium processor's chip?

- A L1 cache
- B L2 cache
- C Integrated cache
- D DRAM cache

Answers

- 1** C. Superscalar design is the idea that the processor has more than one instruction pipeline to process application code. Choice B describes a feature called *dynamic execution*, and choice D describes a clock double chip. *See "Pentium."*
- 2** D. The Pentium II chip comes in a package called the single edge contact (SEC) chip package, which fits into Slot 1. Sockets 5 and 7 are used by original Pentiums, and Socket 1 is used by some 80486 chips. *Review "Pentium II."*
- 3** B,E. AMD's Athlon, Athlon XP, and the Duron chip are placed in Socket A. *Check out "Don't Forget Non-Intel Chips."*
- 4** B. The original Pentium processor was packaged in the pin grid array (PGA) chip type, which fits into either Socket 5 or Socket 7. Socket 5 and Socket 7 are incorrect choices because they are not chip types; they are the names of sockets. Dual inline package (DIP) is an older type of packaging for processor chips, while SEC is the chip packaging used with Pentium II processors. *Peruse "Pentium."*
- 5** B. Virtual memory is one of the features of protected-mode processors. Choice C describes real-mode processors, the opposite of protected-mode processors. *Take a look at "Real-mode versus protected-mode."*
- 6** D. The Pentium III processor contains 32K of L1 cache. *Peek at "Pentium III."*
- 7** C. Registers are storage areas for information to be processed by the processor. The data bus is the pathway to system memory, the address bus controls how much memory can be recognized by the processor, and the math co-processor performs many of the complicated mathematical operations. *Look over "Registers."*
- 8** D. The Pentium Pro has a 36-bit address bus, which enables it to access 64GB of system memory. The original Pentium processor could address up to 4GB of system memory. *Study "Pentium Pro."*
- 9** B. The math co-processor performs many of the complicated math operations on behalf of the CPU, while the CPU itself performs the logic functions. *Refer to "Math co-processor."*
- 10** B, D. The Pentium processor is placed into either Socket 5 or Socket 7. Pentium II uses Slot 1, and some 486 chips use Socket 1. *Examine "Pentium."*
- 11** C. The Pentium II chip increased the L1 cache to 32K, which is broken into two 16K channels. One channel is used for application code and one for data. 486 chips originally used 8K, and Pentium chips originally used 16K. *See "Pentium II."*
- 12** B, C. Pentium 4 processors have either 423 or 478 pins, so they are placed in either Socket 423 or Socket 478. *Review "Pentium 4."*

Picking Your Processor

- 13** D. The Celeron chip is a cut-price Pentium chip. The L2 cache is decreased to save on cost. Celeron chips either have no L2 cache or 128K cache, while the Pentium II, Pentium III, and Pentium 4 chips have 512K of L2 cache. *Check out "Celeron."*
- 14** B. The address bus dictates how much memory the CPU can address. The data bus is the pathway to system memory, and registers are storage areas for data being processed. *Peruse "Address bus."*
- 15** D. The single edge connector (SEC) was used for Pentium II processors. Slot 1 is not a chip package type, but a slot that holds the SEC. Newer processors now use the PGA format. Although DIP chips are no longer used for processors, it is still important to understand what the DIP chip looks like because most BIOS chips come in that format. *Take a look at "Chip packaging."*
- 16** B. The Pentium processor has two instruction pipelines, named U and V. *Peek at "Pentium."*
- 17** C. The PGA packaged chips can be placed in ZIF sockets. These sockets have a lever; when the lever is pulled, the processor chip pops out of the socket. *Look over "Chip packaging."*
- 18** C. The Pentium III processor runs on 133 MHz motherboards. The Pentium and the K6 run on either 60 or 66 MHz motherboards, while the Pentium II can run on 100 MHz motherboards. *Study "Pentium III."*
- 19** B. The Pentium 133 is a clock double processor. It runs twice as fast as the motherboard. Therefore, the motherboard speed is 66 MHz. *Refer to "Pentium."*
- 20** A. L1 cache is integrated into the processor, while L2 cache typically has always resided outside the CPU, usually on the motherboard. DRAM is a type of memory used for RAM. It is not used for cache memory. *Examine "Cache memory."*

Chapter 3: What to Remember about Memory

Exam Objectives

- ✓ Understanding memory terminology
- ✓ Identifying the types of RAM
- ✓ Understanding the types of DRAM
- ✓ Working with memory modules
- ✓ Identifying parity and non-parity memory

Finding out how much memory a computer has is one popular way to measure the computer's power and capabilities. Think about it: If someone asked you what kind of computer you have, what would you say? Probably something like, "I have a Pentium 4 with 1024MB of RAM." But why do we measure the power of a computer based on the amount of memory it has?

In this chapter, you discover the purpose of memory and some of the different types of memory that are found in computers today. This chapter also discusses issues that affect the installation of memory in personal computers and laptops.

Understanding the Types of Memory

This section outlines different types of computer memory. The term *memory* refers to anything that stores information either permanently or temporarily. Computers have two different flavors of memory, ROM and RAM. From an exam perspective, make sure you fully understand the different types of memory and their uses because about 10% of the A+ exam focuses on memory.

Remembering the purpose of memory

Before we look at the different types of memory, let's first ensure that you understand the purpose of memory. We can compare memory to your desk at home or in the office. Whether sitting at your home or office desk (working on a proposal or preparing for your A+ Certification exam), chances are your desk is covered with documents, books, and papers. This desk is your

work area, and its size dictates how many documents you can work on at any given time.

System memory works the same way. You have documents and applications stored on the hard drive. When you want to work on these documents, you open them and place them in the computer's work area. The work area (or desk space) for a computer is system memory. When you want to work with any application or document, the computer must retrieve that information from the hard drive and execute it from memory.

Assume, for instance, that your computer has 512MB of memory (not a lot in this day and age). You start up your system, which is running Windows XP, and decide to run Microsoft Word and Adobe Photoshop at the same time. Assume that you have opened two very large files in each application. Assume further that you are using 480MB of precious memory at this point — a few MBs for the operating system to load, and a few for each running application. As you can see, your memory usage adds up quickly.

In this scenario, you have already used 480MB of memory, which leaves 32MB of memory remaining. Assume that you are about to open up a Photoshop document and copy and paste information from one file to another. To put it simply, you are running out of desk space. You can solve the problem in one of two ways: You can either do less work (in other words, work on one application at a time — although this solution would not serve business users very well because they often need to run multiple applications simultaneously), or you can get a bigger desk, which in computer terminology means *installing more RAM*. When you install more RAM, you have a bigger desk to work on.

Now that you understand the general purpose of memory, let's dive into the different types of memory. There are many different types of memory you are required to know for the A+ exam that are outlined in the next sections.

Read-Only Memory (ROM)

Read-Only Memory (ROM) is a type of memory that you cannot write to. Information is written to ROM chips by the manufacturer, and this information cannot be changed. In the past, if ROM information needed to be updated, you had to remove the original ROM chip and replace it with an updated ROM chip from the manufacturer. Today you can update the ROM by running a special software program downloaded from the manufacturer's Web site, which means that you don't really have a ROM chip — you have an EEPROM (more on EEPROM in a bit).



Software written to a ROM chip is called *firmware*.

One of the major uses for ROM is storing the system *BIOS (Basic Input-Output System)*, which contains *Power-On Self-Test (POST)* routines and other

routines that initiate the loading of the operating system. The BIOS also contains the low-level code that allows the system to communicate with hardware devices.



For the exam, you need to know that the POST is part of the BIOS code stored in ROM. The POST contains routines that initiate the loading of the operating system and routines that make possible the communication between hardware devices.

EPROM

Erasable Programmable Read-Only Memory (EPROM) is a type of memory that normally cannot be written to because it is a variation of ROM. An EPROM chip is a special ROM chip that the manufacturer can reprogram by using a special programming device that uses ultraviolet light.

EEPROM

A new implementation of ROM is called *Electrically Erasable Programmable ROM (EEPROM)*, or *flash ROM*. The manufacturer writes the software instructions into the ROM chip, but you can update these instructions by running a special software setup program provided by the manufacturer. The software setup program is usually provided to you through the manufacturer's Web site.



For the exam, remember that EEPROM, better known as flash ROM, is a ROM chip that can be rewritten with special EEPROM update software provided by the manufacturer of the chip.

EEPROM has become the typical way to update your BIOS. BIOS code is designed to work with certain hardware. As hardware improves, you need to update your BIOS code so that your system is aware of these hardware improvements. Therefore, the manufacturer places BIOS updates on its Web site for computer users running its particular BIOS to download. You just have to download the BIOS update program and then run the BIOS update on your system. The update rewrites the BIOS instructions, making the computer "more aware" of today's hardware.

Random Access Memory (RAM)

Of the two flavors of memory (ROM and RAM), RAM is probably the more fundamental. ROM is permanent memory, or permanent storage of information. As the computer's primary working memory, *RAM*, or *Random Access Memory*, stores information temporarily. RAM is volatile, meaning that it needs constant electrical current to maintain the information that resides in its chips. If the electrical current is lost, the contents of RAM are erased. When the computer is powered off, all the contents of RAM are flushed out.

The following sections discuss the different types of RAM. On the exams, you can expect a few questions about the different types of memory, so be sure that you are familiar with these different types of RAM.

DRAM

Dynamic RAM (DRAM) is probably the most popular type of memory today and the one that you are most often going to upgrade. When someone says to you, “I have 1024MB of Dynamic RAM,” he or she is talking about DRAM.

Dynamic RAM gets its name from the fact that the information stored in DRAM needs to be constantly refreshed. Refreshing involves reading the bits of data stored in DRAM and then rewriting the same information back. DRAM is single ported — meaning that you can read and write to the memory but not at the same time.

Older implementations of RAM measured the memory’s performance based on the time it took the CPU to access that data. The measurement used to determine the speed of memory is *nanoseconds (ns)* — one nanosecond equals a billionth of a second. If you have memory that is 50 ns, and your best friend has memory that is 70 ns, your memory is presumably faster. Your CPU receives the information from memory after waiting only 50 billionths of a second, whereas your best friend’s CPU waits 70 billionths of a second.



The lower the number of nanoseconds, the better the performance.

The speed of older DRAM ranges from 60 ns to 80 ns. Today’s implementations of DRAM measure the speed of memory in megahertz (MHz) — typically matching the motherboard speed. For example, my Pentium II system uses 100 MHz memory because it runs on a 100 MHz motherboard.

For more information on the types of DRAM, see the section, “Identifying the Types of DRAM,” later in this chapter.

SRAM

Static RAM (SRAM) — so-called because the information held in its memory cells doesn’t need to be refreshed — requires less overhead than DRAM to maintain the information stored in memory.

With speeds running from 10 ns to 20 ns, SRAM is much faster than DRAM. Because SRAM is faster memory than DRAM it is also more expensive, which is why people add DRAM to their systems more often than they add SRAM.

SRAM is typically used for cache memory. *Cache memory* stores frequently used data and program code after it is read from slower DRAM. Think of cache memory as a bucket that sits beside the CPU and stores frequently

used information. After the system has searched through DRAM once for specific information, it can store that information in the bucket for easy access later. The next time the data is requested, it is read from cache instead of from system memory.

Because cache memory is much faster than DRAM, the CPU first tries to retrieve the information from cache, specifically L1 cache first and then L2 cache. If the information is not located in cache, the system then tries to retrieve the information from memory. If the information is not located in system memory, it then is retrieved from disk. Attempting to retrieve the requested information from cache first reduces wait time if the information actually resides there because of how fast cache is compared to DRAM.



Cache memory (SRAM) stores frequently used data and program code. Because cache memory is faster than DRAM, retrieving information from cache means that the processor does not have to wait for the slower DRAM, thus enhancing system performance.

CMOS RAM

The *Complementary Metal-Oxide Semiconductor (CMOS)* is the area where the computer stores its configuration information, such as whether or not the computer has a floppy drive, the amount of memory installed, the date and time for the system, and the number and size of the hard drives that are installed. Think of the CMOS information as an inventory list for the majority of components that are installed on the computer. For more information on CMOS, see Book 2, Chapter 4.



CMOS is the computer's inventory list. It tells the computer which devices reside in the system. For example, the CMOS information lists your hard drive, floppy drive, and other information — such as the date and time for the system.

Where is the CMOS information stored? Is the CMOS information stored in the BIOS chip, or perhaps another ROM chip? The answer is neither. In fact, if the information were stored in a ROM chip, you wouldn't be able to go into the CMOS setup program and change the configuration. The CMOS configuration information is stored in a type of RAM called *CMOS RAM*.

CMOS RAM is a special, volatile RAM chip that stores the CMOS information. *Volatile* means that if power is lost, the information is wiped out. This could present a problem with regard to CMOS configuration because if the CMOS RAM is wiped out, the computer forgets its inventory information and has to relearn it. Thus, the computer has a small battery on the motherboard that maintains enough of a charge to avoid CMOS RAM losing its data.

156 *Identifying the Types of DRAM*



For the exams, remember that CMOS information is stored in CMOS RAM, which is volatile memory that maintains its information by using a battery stored on the system board.

Shadow RAM

Part of the boot process involves copying some of the BIOS instructions from ROM up to RAM and then executing those instructions from RAM rather than from the ROM chip. Why? Because ROM is much slower than RAM, performance speed increases when executing the instructions from RAM instead of from ROM. The process in which a copy of the BIOS instructions is shadowed, or copied, to an area of memory called *shadow RAM* is called *shadowing*.

VRAM

Video RAM (VRAM) is dual-ported memory, meaning it can be read from and written to at the same time. DRAM is single-ported, which means that the memory can be written to and read from, but not simultaneously — only one direction at a time. VRAM, however, lets you do both simultaneously.

VRAM is most commonly used on video accelerator cards to store the values of the pixels on the screen for refresh purposes. VRAM is the favored memory for video because it outperforms the other memory types by being dual ported.

WRAM

Window RAM (WRAM), also known as *Window Accelerator Card RAM*, is a modification of VRAM and is also used for video display purposes. Like VRAM, WRAM is dual-ported memory but runs about 25 percent faster. In general, WRAM offers better performance than VRAM.

Identifying the Types of DRAM

Dynamic RAM (DRAM) is the most popular type of memory used in systems today. It is also the most popular type of memory that computer users add to their computers for the purpose of upgrading memory. Therefore, you must understand the different types of DRAM and what types of DRAM outperform others.

Standard DRAM

Memory is organized into rows and columns, like a spreadsheet. The information is stored in the different cells, or blocks, that are created by the intersection of these rows and columns. With *standard DRAM*, the CPU requests data by sending the address of the row and the address of the column for every

block of data that needs to be read to the memory controller. The memory controller then fetches the information from that memory location. Figure 3-1 shows two memory cells that hold data that the CPU wants to have.

	0	1	2	3	4	5
0						
1						
2						
3						
4						

Figure 3-1: Looking at how data is accessed in memory.

To access the information shown in Figure 3-1, the CPU follows these basic steps to request information from standard DRAM:

1. In the first clock cycle, it sends the row address (1).
2. In the second clock cycle, it sends the column address (2).
3. On the third clock cycle, the memory controller reads the information (Address 1-2).
4. In the fourth clock cycle, the row address for the second memory cell is given (1).
5. In the fifth clock cycle, the column address for the second memory cell is given (4).
6. In the sixth clock cycle, the second memory cell is read (Address 1-4).

Fast page mode

Fast Page Mode (FPM) improves the performance of standard DRAM by not requiring a row address for each request to memory, assuming that the next block of data is on the same row (which in most cases will be true). The following list outlines the basic steps to access the same two blocks of data shown in Figure 3-1 via fast page mode memory:

1. In the first clock cycle, the CPU sends the row address (1).
2. In the second clock cycle, it sends the column address (2).
3. On the third clock cycle, the memory controller reads the information (Address 1-2).
4. In the fourth clock cycle, the column address is given (4).
5. In the fifth clock cycle, the second cell address is read (Address 1-4).

You can see in this example that it takes less time to read both blocks of data from memory with fast page mode DRAM. Therefore, FPM memory is a faster DRAM memory type than standard.

Extended data output

Extended Data Output (EDO) memory is about 10 to 15 percent faster than FPM memory and is usually found on 66 MHz motherboards. With EDO memory, the memory controller can read data from a memory block while listening for the next instruction. This capability increases performance because the memory controller doesn't have to wait for the next instruction after reading a block of memory; while it is reading one block of memory, it is receiving the next instruction. In contrast, with FPM DRAM, reading one memory block and listening for the next instruction are done in multiple steps.

Burst Extended Data Output

Burst Extended Data Output (BEDO) is a bursting-type technology. The word *burst* refers to the fact that when one memory address is requested and that address is retrieved, the system bursts into the next couple of blocks and reads those as well. The theory behind BEDO is that the system has already gone through the trouble of locating that block, and chances are that the next request will be for the next block, so why not take that information while the memory controller is already there? If that extra block is the next requested block from the CPU, the memory controller already has the data and can pass it to the CPU immediately.

BEDO is 50 percent faster than EDO. Because of lack of support from computer manufacturers, however, BEDO has not been used in many systems. It has been surpassed by SDRAM instead.

Synchronous DRAM

Synchronous DRAM (SDRAM) is memory synchronized to the system board speed. This synchronized speed means that the data stored in memory is refreshed at the system speed, and data is accessed in memory at the system speed as well.

SDRAM is one of the most popular types of DRAM found in later Pentium systems, such as the Pentium II. When you upgrade memory on your system, if you determine that you need SDRAM, you will then need to determine what speed SDRAM. Because you are running at the system speed, you must match the DRAM speed with the motherboard speed. Thus, if you have a 100 MHz motherboard, you need 100 MHz SDRAM. If you have a 133 MHz motherboard, you need 133 MHz SDRAM.



As mentioned, if you have a 100 MHz motherboard, you will purchase 100 MHz memory, typically labeled PC100. Be aware, however, that there is some flexibility when purchasing SDRAM. For example, I have a 100 MHz motherboard on an old Pentium II system. When I upgraded the DRAM on this system, I couldn't buy PC100 memory because PC133 (which is SDRAM that runs at 133 MHz) was the popular memory at that time. Not a problem! You can use faster memory than your motherboard speed as long as you are willing to accept that you have paid for memory that will not run to its full potential speed. In my example, the 133 MHz memory is only running at 100 MHz due to the speed of the motherboard.

Rambus DRAM

At the time that SDRAM was popular, there was a high-speed flavor of DRAM on the market called *Rambus DRAM (RDRAM)*, which runs at speeds around 800 MHz! The RDRAM chips have a 16-bit internal bus width and are packaged together in a 184-pin, gold-plated memory module called a *Rambus Inline Memory Module (RIMM)*. In order to take advantage of this type of memory, you need a motherboard and chipset that support RDRAM.

DDR

Double Data Rate (DDR) memory gets its name from the fact that it can transfer data twice during each clock cycle, whereas SDRAM can transfer data only once per clock cycle. DDR memory ships in 184-pin DIMM modules (see the section "DIMMs," later in this chapter) for desktop computers and 200-pin SO-DIMMs for laptop systems.

The speed of DDR memory is measured in MHz, like SDRAM is, and is labeled to indicate the speed. The labeling of DDR memory may look obscure at first because it also indicates the bandwidth by taking the speed and multiplying it by 8 bytes of data (64 bits). So if DDR memory is labeled PC1600, that label breaks down like this: If you divide the 1600 by 8 bytes, you get the speed of the memory; in this case, you're looking at 200 MHz memory. PC2700 runs at 333 MHz, while PC3200 runs at 400MHz. When you upgrade memory on systems that require DDR memory, you need to know the speed of the DDR memory.

DDR2

Improvements to DDR memory have already started with DDR2 memory. *DDR2* memory runs at speeds 400 MHz and higher, which is where DDR memory left off. DDR2 memory uses 240-pin memory modules and runs at 1.8 volts (as opposed to 2.5 volts for DDR memory). This results in less power consumption for more memory — which is great for laptop users.

Popular modules of DDR2 memory at the time of this writing are PC3200 (400 MHz), PC4200 (533 MHz), PC5300 (666 MHz), and PC6400 (800 MHz).

How Would You Like Your Chips Packaged?

Whether you're purchasing or installing RAM, understanding the different types of memory packages available is important. The following sections identify different memory packages used in desktop computers and laptop systems.

SIMMs

Single Inline Memory Modules (SIMMs) used to be one of the most popular types of memory modules, but they have been replaced by DIMMs (see the next section) in recent years. A *SIMM* is a card that holds a number of memory chips and has an edge connector containing a number of pins that make contact with the motherboard. This design makes it quite a bit easier to install memory than it was many years ago. In the past, you had to take a dual inline package (DIP) chip out of the system board and reinsert a new chip. Today, you purchase a card of chips (a SIMM) and install the SIMM into one of the SIMM sockets.

SIMMs come in two flavors, 30-pin and 72-pin, which describe the number of connectors that make contact with the motherboard. Before buying a SIMM to install in a computer, review the documentation for the computer or look at the system board to determine what size SIMM module you need. Figure 3-2 shows a 30-pin SIMM, a 72-pin SIMM, and a 168-pin *dual inline memory module (DIMM)*.

The 30-pin SIMMs have an 8-bit data path, meaning they supply information in 8-bit blocks. When installing memory into a system, you *must* install enough SIMMs to fill a memory bank. A *memory bank* is the number of SIMMs it takes to fill the data path of the processor. For example, if you have a system with a 486 processor, the processor is a 32-bit processor. Therefore, the processor wants to deal with information in 32-bit chunks. When using 30-pin SIMMs, you need to install four of them at a time to fill a memory bank because each 30-pin SIMM only supplies 8 bits of data ($8 \text{ bits} \times 4 \text{ SIMMs} = 32\text{-bit chunks}$).

The 72-pin SIMMs supply information in 32-bit chunks. Therefore, if you are installing 72-pin SIMMs on a system using a 32-bit 486 chip, you need just one SIMM to fill a memory bank and the data path. If you're installing 72-pin SIMMs in a Pentium system, you must install SIMMs in pairs because the Pentium data path is 64-bit; to fill a bank on these systems, you need two 32-bit modules (72-pin SIMMs).

For the exam, remember the data path of the SIMM modules. You should also know how many SIMMs it takes to fill a memory bank on different systems. Remember that a *memory bank* is the number of memory slots needed to fill the data path of the processor.






Figure 3-2: Looking at SIMM and DIMM memory modules.

In this day and age, you most likely will not see SIMMs in a system unless you are supporting older computers. If you see a system that uses SIMMs, it most likely conforms to the 72-pin format.



You can easily distinguish what size SIMM a system uses, even if you don't have the documentation for that system. The 72-pin SIMMs have a notch close to the center of the module. If there are SIMMs already installed in the system, you can take them out and examine them. They usually have a label with a 1 or a 72, representing the pin numbers, at either end of the module — so if you see a number 72, you know you have a 72-pin SIMM.

DIMMs

Dual Inline Memory Modules (DIMMs) are like SIMMs, only they supply information in 64-bit chunks. DIMMs use 168 pins on the module and are a little larger than the 72-pin SIMMs. (Refer to Figure 3-2.)



For the exam, remember that SIMMs come in 30-pin and 72-pin flavors, whereas DIMMs have 168 pins. DIMMs are also the most popular type of memory module that you will find in systems today.

Consider the memory bank issue again. Because the DIMM supplies data in 64-bit chunks, and the data path of a Pentium processor is 64-bit, you can install DIMMs singly in a Pentium system. On the other hand, you must install SIMMs in pairs in a Pentium system. Figure 3-3 shows what 72-pin SIMM and 168-pin DIMM sockets look like.

SODIMM

Small Outline Dual Inline Memory Modules (SODIMMs) are memory modules that are smaller than normal DIMMs and are used in laptops. A SODIMM comes in three different-sized modules: a 32-bit 72-pin module; a 64-bit 144-pin module (SDRAM); and a 64-bit 200-pin module (DDR). Figure 3-4 compares a SODIMM and a DIMM.

MicroDIMM

A *Micro Dual Inline Memory Module (MicroDIMM)* is another memory module that is used in laptop computers. The MicroDIMM is smaller than the SODIMM and comes in a 144-pin module for SDRAM and a 172-pin module for DDR memory.




Figure 3-3: Looking at memory sockets on a motherboard.




Figure 3-4: Comparing a SODIMM and a DIMM.

Understanding Error-Checking Memory

There are two primary types of error-checking memory that have been used in systems over the years. The following sections introduce you to these two types of error-checking memory — be sure to become familiar with them for the exam.

Parity versus non-parity

In this section you learn about parity versus non-parity memory. *Parity memory* is a type of error-checking memory, which is memory that verifies the information stored in memory is what is actually read from memory at a later time. *Non-parity memory* is simply memory that does not perform any kind of error checking to ensure that the data written to memory is what is actually read when it is retrieved. Let's look at how parity memory works!

There are two types of parity memory: *odd parity* and *even parity*. Both parity methods function the same way but differ in the sense of whether they look for an odd number of bits or an even number of bits. This discussion uses odd parity as the example.

With parity memory, for every byte (8 bits) of data written to memory, there is an additional 9th bit known as the parity bit. When storing information to memory, the number of the enabled data bits (bits set to 1) written to memory are added up.

With *odd parity*, if an even number of data bits are enabled, the parity bit is set to 1 (enabled) so that there is an odd number of enabled bits in total written to memory. If the result of all the enabled data bits is odd, the parity bit is set to 0 (disabled) so that the odd number of enabled bits is retained.

After the parity bit has been set, the byte of data and the parity bit are written to memory. Note that *even parity* works the same way, only it looks for an even number of enabled bits; if the number of enabled bits is odd, then the parity bit is enabled.

When the CPU requests data from memory, the data byte is retrieved along with the parity bit that was generated when the byte of information was stored in memory. The system looks at the data byte and calculates whether the parity bit stored in memory should be set to 1 or 0. It then compares the answer it has just generated with the value of the parity bit stored in memory. If the two match, the integrity of the information in memory is considered okay, the parity bit is stripped from the data byte, and the data is delivered to the CPU. If the two differ, you have a *parity error*, meaning that there is a problem with the integrity of the data stored in memory.



Note that parity memory cannot correct the error; it just reports that an error exists.



For the exam, remember that parity memory has an extra bit (the parity bit) for every 8 bits of data. SIMMs with parity come in 9-bit (30-pin SIMM) or 36-bit (72-pin SIMM) flavors. Also, remember that a parity error indicates that there's something wrong with the integrity of data stored in memory.

ECC memory

Error-checking and correction (ECC) memory is memory that can detect data integrity problems the way that parity memory can, the difference being that ECC memory can recover from the error and attempt to fix the problem with the data being read.

Working with Cache Memory

Cache memory stores frequently used data and program code after it is read from slower DRAM. Cache memory is made up of SRAM, which is much faster than DRAM. The average speed of DRAM is 60 ns, whereas the average speed of SRAM is 20 ns. If at all possible, you want the CPU's request for information to be serviced by cache memory for a quicker response. To help service these responses, the system has two major levels of cache memory: L1 and L2 that are popular, and also an L3 cache that is making its way on systems today.

L1 cache

Level 1 cache, or *L1 cache*, is “internal cache” integrated into the CPU. This memory is typically a small amount of SRAM integrated into the processor’s chip, giving the processor instant access to this memory with no wait time. *Wait time* is the amount of time it takes between the processor requesting information stored in memory and actually receiving the information.

Every processor before the Pentium processor has L1 cache integrated into the processor chip, but the amount of L1 cache can vary. For example, the 486 chips had 8K of L1 cache, whereas the early Pentium processors had 16K of L1 cache. Newer processors have doubled that amount to 32K of L1 cache.

L2 cache

Level 2 cache, also known as *L2 cache*, exists outside the CPU, usually on the motherboard or just outside the processor but in the processor casing. Therefore, some delay occurs when the processor accesses the information in L2 cache due to the distance between the processor and the L2 cache.




For the exam, remember that L1 cache is SRAM integrated into the processor’s chip, whereas L2 cache is SRAM located outside the CPU, usually on the system board or in the casing of the processor.

One of the selling points of different processors is the amount of cache memory that comes with the processor. Many processors today typically have at least 32K of L1 cache and 512K, 1MB, or 2MB of L2 cache inside the casing of the processor. The more cache memory a system has, the bigger the bucket to store more frequently used information.

When the processor retrieves information, it first checks to see whether the information it needs is stored in L1 cache (because L1 cache has no wait time). If the processor does not find the information in L1 cache, it checks the L2 cache. If the information cannot be found in either L1 or L2 cache, the processor finally retrieves the information from RAM. Figure 3-5 shows the steps the processor takes to retrieve information.

Figure 3-5:
How the CPU
retrieves
information
from
memory.



L3 cache

Because processors today provide a small amount of L1 cache *and* a large amount of L2 cache, some people are now using the term *L3 cache* to identify cache that resides on the motherboard.

If you are confused by L3 cache I can see why — we have changed the terminology on you. In the past, L2 cache resided on the motherboard but now that processors include L2 cache on the CPU, the term for cache memory located on the motherboard is L3 cache!

Installing or Upgrading Memory

The discussion in this section focuses on issues related to memory upgrades. In general, upgrading memory is a simple task — assuming you purchase the proper type of memory for the upgrade. Factors that affect the proper type of memory are

- ◆ Type of memory (FPM, EDO, SDRAM, DDR)
- ◆ The speed of the memory
- ◆ Pin connector type
- ◆ Parity versus non-parity

The following sections discuss each of the factors that affect how you upgrade your system's memory.

Type of memory

The first thing you need to know to upgrade your computer's memory is which type of memory you need. You first need to figure out whether you need to install a SIMM, DIMM, or SODIMM. Pentium II and later desktop computers usually need a DIMM; laptops use a SODIMM.

After you have determined the memory module type, you need to determine the type of memory to install, such as SDRAM or DDR memory.

Speed

When you buy memory, you need to take into account the speed of the memory. Older memory, such as FPM or EDO memory, is typically measured in nanoseconds (ns) and ranges from 60 ns to 80 ns. With these types of memory, it is important not to mix speeds or the system will become unreliable. The speed of a SIMM is usually indicated on the chips themselves

(displaying either a numeric value or a simple minus sign with a number). For example, a memory module running at 70 ns would show either “70” or “-7” on the chips.

The speed of newer memory types, such as DIMMs and SODIMMs, is measured in MHz. If you’re buying SDRAM or DDR memory, make sure you verify that you get the correct memory speed. For example, I recently upgraded the memory in the laptop I’m using to write this book, and when I went to the store to buy a memory module, the first questions the in-store computer geeks asked were, “What type of memory?” and “What speed?” The system documentation can help you determine the speed of memory needed. If you don’t have the documentation, be sure to look up the information on the manufacturer’s Web site.

Connectors

Another important issue with regard to memory installation is with the metal used on the memory modules. You need to purchase memory modules that use pins plated with the same metal used in the memory socket on the motherboard. Memory modules use silver or gold plated pins. If the SIMM socket, for example, uses silver-plated connectors, the memory module you purchase must use silver-plated pins. If you mix metal types, you’ll eventually have an unstable system.



Today’s preferred memory modules are gold-plated DIMM modules. It is best not to assume you are using gold plated DIMMs, but to ensure you are installing the correct type of memory by checking the documentation for your system.

Parity versus non-parity

The final issue with regard to memory upgrades is whether the system uses parity or non-parity memory. This information can be determined by checking the documentation that came with the system or by checking the system summary in CMOS. If you can’t find the information in CMOS or have misplaced the documentation, you can try to find the information on the Internet at the vendor’s Web site. When you locate the information, use Table 3-1 to record the type of memory your system has so that you may refer to this when you perform a memory upgrade.

Table 3-1 Identifying Memory Used by Your System

Memory Type (SDRAM, DDR)
Memory Speed (60 ns, 100MHz)
Gold or Silver pins
SIMM or DIMM

Installing memory on desktop PCs

Now that you have purchased the correct type of memory — meaning the correct type at the correct speed with the correct number of pins — you are ready to install the memory!

Take off the computer's cover. You should see either SIMM sockets or DIMM sockets — maybe even both types on an older system.



Look at the sockets and determine whether the memory modules will sit diagonally or vertically. This step is very important. I have seen many people struggle to install memory because they didn't understand how to correctly place the modules in the sockets.

If the socket is on a diagonal, lightly place the memory module vertically and then lay it back diagonally; it should just snap in. If the socket is vertical, place the memory module at a 45-degree angle and lightly lay it back to the vertical position; it should snap in. Figure 3-6 shows the installation of memory modules.




Figure 3-6:
Installing
memory
modules on
a system
board.

When you install the memory module, make sure you line up Pin 1 on the memory module with Pin 1 in the socket. To locate Pin 1 on a SIMM, look for the cutout on the memory module (shown in Figure 3-7) and place it over the shoulder of the SIMM socket.

To locate Pin 1 on a DIMM, simply look at the memory module; the pins are labeled. The DIMM cannot be placed in backward because it's *keyed* (meaning that it has cutouts so that the memory can only be inserted one way). When installing the DIMMs, you will push down firmly on the module to place the DIMM deep in the socket.




Figure 3-7:
Locating
Pin 1 on a
SIMM
memory
module.



Lab 3-1 and lab 3-2 can be found in the `Labs.pdf` file in the `Author` directory of the CD-ROM. Lab 3-1 will allow you to practice finding out what type of memory you have, while Lab 3-2 gives you practice installing the memory on a desktop system.

Installing memory on laptop systems

Installing memory on today's laptop systems is just as easy as installing memory on a desktop PC. The first step is to make sure you buy the correct type of memory for the laptop — again, check the laptop documentation before heading to the store!

To install your newly purchased memory, flip the laptop over. You'll notice a door on the bottom of the laptop that can be removed (shown in Figure 3-8). This is where you add memory.

After the cover has been removed, you can insert the SODIMM into an empty slot by lightly placing the SODIMM into the memory slot on a slight angle and then clamping it back into place, as shown in Figure 3-9. If you are replacing memory, you will need to first remove the old memory module by pressing on the clips on the side and then lifting the memory out.

When you have the SODIMM locked in place, put the cover back on and you're ready to use your laptop with the new memory!



Figure 3-8:
Removing
the cover to
add memory
to a laptop.




Figure 3-9:
Inserting the
SODIMM
module into
the memory
slot.



Lab 3-3 lets you practice finding out what type of memory is needed by a specific laptop. In this lab you will need to research the laptop on the Internet. Lab 3-3 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Getting an A+

This chapter provides an overview of the different types of memory and installation of memory. The following points are touched on:

- ◆ *Read-Only Memory (ROM)* is memory that can be read from but not written to, and the information stored there is permanent.
- ◆ *Random Access Memory (RAM)* is volatile memory that can be written to and read from; information stored there is flushed out when power is lost.
- ◆ *Dynamic RAM (DRAM)* is memory that needs constant refreshing of its memory cells; it is also the type of memory that is typically upgraded on systems.
- ◆ *Static RAM (SRAM)* is static memory, meaning that it does not need refreshing as often as DRAM does. SRAM is also faster than DRAM. SRAM is typically used for cache memory.
- ◆ *L1 cache* is cache memory integrated into the processor, whereas *L2 cache* is cache memory that exists outside the processor's chip.
- ◆ Memory is installed with *memory modules*. 72-pin SIMMs were popular in original Pentium systems, but 168-pin DIMMs are now the popular memory module for desktop systems. Laptops use SODIMMs as their memory module type.
- ◆ You need to know the *type of memory* your system uses before you upgrade your system. For example, you need to note whether your system uses SDRAM or DDR memory. Also be sure to note the speed of the memory.

Prep Test

- 1** Which of the following types of memory is used for cache memory?
 - A SRAM
 - B DRAM
 - C EDO RAM
 - D SDRAM

- 2** How many bits make up the data path of a 30-pin SIMM?
 - A 12
 - B 8
 - C 32
 - D 64

- 3** Which type of cache memory resides on the system board?
 - A L1 Cache
 - B DDR
 - C L2 Cache
 - D SDRAM

- 4** Which type of memory must be constantly refreshed?
 - A DRAM
 - B SRAM
 - C VRAM
 - D L1 Cache

- 5** Which type of memory is dual-ported (can read from and written to at the same time)?
 - A VRAM
 - B DRAM
 - C SRAM
 - D L1 Cache

- 6** Of the following, which is the fastest memory?
 - A 30 ns
 - B 10 ns
 - C 120 ns
 - D 70 ns

- 7** How many 72-pin SIMMs make up a memory bank on a Pentium system?
- A 3
 - B 1
 - C 2
 - D 4
- 8** How many pins does a DIMM have?
- A 64
 - B 128
 - C 168
 - D 32
- 9** How many bits of parity information are found in a 72-pin SIMM?
- A 1
 - B 2
 - C 3
 - D 4
- 10** If memory data is corrupted, what type of error is created?
- A File access error
 - B Parity error
 - C Disk error
 - D DRAM error
- 11** When installing SIMM memory, what do you look for to identify Pin 1 on the memory module?
- A Square solders
 - B Red ribbon wire
 - C Blue ribbon wire
 - D SIMM cutout
- 12** Which of the following best describes VRAM?
- A Memory that has video instructions burned to it by the manufacturer and that cannot be changed
 - B Dual-ported memory usually found on video accelerator cards
 - C Memory stored in the CPU and used for storing frequently used instructions
 - D Memory that has video instructions burned to it by the manufacturer and that can be reprogrammed only by using special software

What to Remember about Memory

- 13** What type of memory is typically used for storing computer BIOS code?
- A DRAM
 - B SRAM
 - C RAM
 - D ROM
- 14** Which of the following best describes shadow RAM?
- A Shadow RAM is an area of system memory where some of the BIOS functions are copied during the boot process.
 - B Shadow RAM is an area of system memory that stores the information twice so that if the system crashes there will be a backup copy.
 - C Shadow RAM is the process of having frequently used data stored in cache for quicker access.
 - D Shadow RAM is the same thing as SRAM.
- 15** How many 30-pin SIMMs are required to fill a memory bank on an 80486 processor?
- A 1
 - B 2
 - C 3
 - D 4
- 16** Which type of DRAM is synchronized with the system clock?
- A EDO
 - B BEDO
 - C FPM
 - D SDRAM
- 17** How many DIMMs are used to fill a memory bank in a Pentium III system?
- A 1
 - B 2
 - C 3
 - D 4
- 18** When a CPU needs to access data, where will it look for this data first?
- A L2 Cache
 - B RAM
 - C L1 Cache
 - D ROM

19 Which of the following best describes odd parity memory?

- A When a byte of data is stored, the values of all of the bits in the byte are added up; if the number is even, the parity bit is enabled to create an odd number. If all of the bits in the data byte add up to an odd number, the parity bit is turned off, leaving the total an odd number.
- B All of the data stored in memory is added up, the answer is then stored in the parity bit.
- C When a byte of data is stored, the values of all of the bits in the byte are added up; if the number is odd, the parity bit is enabled to create an even number. If all of the bits in the data byte add up to an even number, the parity bit is turned off, leaving the total an even number.
- D Every 8 bits of data are added up and the total is stored in the parity bit for that particular 8 bits of data.

20 What is the purpose of CMOS RAM?

- A On system boot, the contents of the BIOS are copied to faster CMOS RAM.
- B CMOS RAM is memory used to store the CMOS information. CMOS RAM maintains its information by using a battery found on the motherboard.
- C CMOS RAM is the amount of memory that is registered into CMOS.
- D CMOS RAM is read-only memory that stores CMOS information.

Answers

- 1 A. SRAM is used for cache memory. DRAM is what is known as system memory and stores data and program code that are currently running. EDO and SDRAM are types of DRAM. *See "SRAM."*
- 2 B. A 30-pin SIMM has an 8-bit data path. 72-pin SIMMs have a 32-bit data path, and a DIMM has a 64-bit data path. *Review "SIMMs."*
- 3 C. L2 cache memory is external to the CPU. This is the memory that is inserted onto the system board as additional cache memory. L1 cache is cache memory integrated into the CPU. *Check out "L2 cache."*
- 4 A. DRAM must be constantly refreshed. The other types of memory do have to be refreshed, but not as often as DRAM. *Peruse "DRAM."*
- 5 A. VRAM, or Video RAM, is dual-ported memory that can read from and written to at the same time. Because of the performance increase with this type of memory, it is the desired memory for video accelerator cards. The other types of memory are single ported, which means that you have to read and write through the same port and can only go one way at a time. As a result, the memory is slower than VRAM. *Take a look at "VRAM."*
- 6 B. With nanoseconds, the lower the number, the faster. Because it is measured in billionths of a second and determines how long it takes the system to access the information stored in memory, 10 ns is faster than the other choices. *Peek at "DRAM."*
- 7 C. 72-pin SIMMs are 32-bit modules. Because Pentium processors have a 64-bit data path, you need two 72-pin SIMMs to fill the data path. *Look over "SIMMs."*
- 8 C. There are 168 pins on a DIMM. The other answer choices for this question are just common numbers used in the computer field, but there are no memory modules that use those numbers of pins. *Study "DIMMs."*
- 9 D. There is always 1 parity bit for every 8 bits of data. Because a 72-pin SIMM is a 32-bit module, it is made up of four 8-bit chunks, so it will have 4 parity bits. *Refer to "Parity versus non-parity."*
- 10 B. You'll receive a parity error if there are any problems with data stored in memory. *Examine "Parity versus non-parity."*
- 11 D. When looking at the memory module, you will notice that one side of the module is cut out. This indicates where Pin 1 is located on the module. There are no ribbon wires on a memory module. *See "Installing memory on desktop PCs."*

- 12** B. VRAM is dual-ported video RAM, meaning that the memory can be written to on one port and read from on another port that is usually found on video accelerator cards. *Review "VRAM."*
- 13** D. A computer manufacturer's BIOS code is typically stored on read-only memory. Today's version of the BIOS chip can be reprogrammed by running special software to update the BIOS. *Check out "Read-only memory (ROM)."*
- 14** A. The term *shadowing* implies "making a copy" — more specifically, copying the BIOS instructions to system memory so that the functions can be executed more quickly. *Peruse "Shadow RAM."*
- 15** D. It takes four 30-pin SIMMs to fill a bank in any 486 processor because 30-pin SIMMs are only 8 bits wide, and a 486 processor's data path is 32-bit. A memory bank's data path should be the same as a processor's data path. If the SIMMs had 72 pins, you'd need only one module because 72-pin SIMMs have a 32-bit data path, the same as the 486 chip. *Take a look at "L1 cache."*
- 16** D. SDRAM memory is synchronized with the system clock speed. If the motherboard runs at 100 MHz, you must buy the same speed memory for the board. *Peek at "Synchronous DRAM."*
- 17** A. Pentium IIIs have 64-bit data paths, and a DIMM has a 64-bit data path as well. So you need only one DIMM chip to fill a bank in a Pentium system. *Look over "DIMMs."*
- 18** C. The CPU checks the L1 cache first because it is SRAM memory integrated into the chip and has no wait time. If the data is not in the L1 cache, the system checks the L2 cache, which is cache memory located outside of the chip, usually on the motherboard. If the data isn't in the L2 cache, then the CPU checks in RAM. If the information is not found in RAM it is then read from disk. *Study "SRAM."*
- 19** A. Odd parity memory stores a parity bit for every byte of data. The parity bit is enabled (set to 1) if the data bits add up to an even number, and is disabled (set to 0) if the data bits add up to an odd number. *Refer to "Parity versus non-parity."*
- 20** B. CMOS RAM is a small amount of memory that maintains the CMOS configuration information. Since this information is stored in a type of RAM, the memory has to have constant power; this power comes from the battery found on the system board. Choice A describes shadow RAM. *Examine "CMOS RAM."*

Chapter 4: Telling Your BIOS from Your CMOS

Exam Objectives

- ✓ Identifying the purpose of BIOS
- ✓ Upgrading a system BIOS
- ✓ Identifying the purpose of CMOS
- ✓ Identifying common settings contained in CMOS

A big part of system configuration is understanding what CMOS is and how the different settings within CMOS can help you troubleshoot a system. A wonderful talent to have as an A+ Certified Professional is the ability to understand CMOS and its functionality. Although CMOS doesn't seem like a place where you'd need to go every day, many PC-related problems can be solved in CMOS.

In this chapter, I introduce you to some of the common CMOS options that are found in modern systems. Keep in mind that these options may be labeled differently in different systems or may not even exist in some systems; the CMOS setup program is unique to each system. For example, the CMOS setup program on a Compaq system differs from that of a Dell — although the concept of CMOS remains the same.

I also make you aware of popular CMOS settings so that when a scenario arises that calls for such a change, you are comfortable with not only what you need to do, but also how to do it.

The BIOS and Its Purpose

I can't start a discussion on CMOS without first comparing it to the system BIOS — the two are closely related and are often confused by many IT professionals.



BIOS stands for *Basic Input-Output System* and is the low-level instructions used to communicate with the system devices. This is different than CMOS, which is simply an *inventory list* for the system, containing information such as the hard drive space and amount of memory that exist on the system. So,

CMOS is the inventory list while BIOS is the actual code that is run to communicate with those devices.

Originally, the BIOS was stored in a ROM chip on the motherboard but because you had to replace the chip with a new chip to update the BIOS, today's systems have the BIOS code stored in an EEPROM chip located on the motherboard (for more information on ROM and EEPROM, check out the memory topic in Chapter 3 of this minibook). The chip that contains the BIOS code is known as the BIOS chip. The BIOS chip should be easy to find — it's usually rectangular and clearly labeled, normally with the name of the manufacturer of the BIOS chip and the date that the chip was created. In Figure 4-1, the BIOS chip is the chip with the white label indicated with a number 1.




Figure 4-1: The BIOS chip is an EEPROM chip located on the motherboard and contains the BIOS code.

Along with the low-level code, the BIOS also contains the POST (Power-On Self-Test) and the CMOS setup program. The POST is a self-diagnostics routine that the system goes through each time it boots up. This self-diagnostics checks to see that each device listed in CMOS actually exists on the system. The POST also tests devices, such as I/O ports and memory, to verify that they not only exist but also that they still function.

You use the CMOS setup program, stored in the BIOS, to navigate through the CMOS settings — but the BIOS does not contain the settings themselves. As you find out in the section, “Understanding CMOS,” later in this chapter, the CMOS settings are maintained in CMOS RAM.

Upgrading the System BIOS

When your system was designed, its BIOS program code was designed to work with very specific devices. As you know, computer technology changes very quickly, almost overnight. So, what can you do if you have an older system and you want to update its capabilities?

The BIOS dictates a system's capabilities. For example, assume that the BIOS on my old Pentium is only aware of an 8GB drive, but I would like to install a 12GB drive into the system. I must first upgrade the BIOS so that the system knows how to refer to a drive of that size. As far as the current BIOS is concerned, there is no such thing as a drive larger than 8GB — I need an upgraded BIOS to make the system understand!

Another example of a good time to upgrade the BIOS code in hardware is with a wireless home router I purchased a few years back. I knew that I wanted to configure the wireless network to limit which MAC addresses (network cards) could connect to the wireless network. Unfortunately, there was no such setting in the configuration screens of the wireless router — so I went to the manufacturer's Web site and updated the BIOS in the router. After I completed the upgrade of the router, the setting I was looking for suddenly appeared on the configuration screens! The point is that it is extremely common to upgrade the BIOS on devices to make sure that the device is up to date with the current trends.

Performing the BIOS upgrade

In the past, with older systems, you would upgrade the BIOS by completely replacing the BIOS chip. Back then, BIOS was stored on a ROM (Read-Only Memory) chip, which could not be written to. In order to “rewrite” the code, you needed to replace the entire chip!

Today's systems use a modified version of the ROM chip — an EEPROM (Electrically Erasable Programmable Read-Only Memory) chip. To upgrade the program code on an EEPROM chip, you don't need to physically replace the chip; you just run a software program that was designed to rewrite the program code.

To upgrade the BIOS, you need to get the update program from the manufacturer. You can usually find the program on the manufacturer's Web site (and it's usually downloadable), or you may be able to order a CD from the manufacturer. Be sure to follow the manufacturer's directions on how to apply the update to your BIOS.

For example, a popular technique of updating a BIOS is to download the BIOS update program from the manufacturer's Web site to a floppy disk. After you download the BIOS program to a floppy disk, you boot off the floppy disk and the update starts. Just follow the directions on the screen. The update program rewrites the program code stored in the BIOS chip that is normally read-only. Because you are writing to this ROM chip with a special program, they call the ROM a *flash ROM*, where *flashing* is the process of rewriting the program code.



Today's computer systems are no longer shipping with built-in floppy drives so it is up to the manufacturer of the BIOS to decide how the BIOS update will be performed. If the manufacturer still requires you to use a floppy disk then you can purchase a USB floppy drive to connect to a system when you do the update. Some manufacturers may allow you to perform a BIOS update by running the update program from CD-ROM after downloading it.

Potential issues with BIOS upgrading

If you decide to perform a BIOS upgrade, consider these few words of caution:

- ◆ **Be sure you're applying the correct BIOS update for system.** You want to make sure that you are aware of who the manufacturer is for your BIOS before you look for an update. Once you know the manufacturer of your BIOS you can go to the manufacturer's Web site and download the update. You will not be able to mix BIOS codes between manufacturers — for example, you cannot update your Compaq system with a BIOS update from Dell.
- ◆ **Be sure that you not only have the BIOS update for the correct system, but also that the BIOS update you have is designed for your version of the BIOS.**

Each BIOS has a revision number or version number. The developers of the BIOS update program may place a version check into the BIOS update, but you cannot be sure of that. So as a result, be sure to double-check that the BIOS update you are running is for the *BIOS version* you have.



Installing the wrong BIOS version can cause the system to become dysfunctional. If you can't find a BIOS version number, use the BIOS date to decide which update to download — either one should be located within CMOS or displayed with the BIOS manufacturer on the screen during boot-up, as shown in Figure 4-2.

- ◆ **When you start a BIOS update, be sure that you complete it.** Rewriting only a portion of the BIOS code (whether as a result of an accidental shutdown or power loss during the writing process) can cause the system to become dysfunctional. The actual BIOS update should take only a few seconds, so try not to disrupt the BIOS update after it starts.

```

Phoenix BIOS 4.0 Release 6.0
Copyright 1985-1998 Phoenix Technologies Ltd. All Rights Reserved
Copyright 1996-1998 Intel Corporation.
404CL0X0.15A.0306.P02

Intel Celeron(tm) processor 333 MHz
128MB System RAM

Legacy Keyboard ... Detected
Legacy Mouse ..... Detected

Fixed Disk 0: QUANTUM FIREBALL EX10.2A-(PM)
ATAPI CD-ROM: MATSHITA CR-588-(SM)
ATAPI Removable Drive: IOMEGA ZIP 100-(SS) ATAPI

Press <F1> to enter SETUP

```

Figure 4-2:
Identifying
the BIOS
version by
starting up
the system.



As a way to prevent accidental writing to the BIOS code, some motherboards ship with a jumper on the board that must be removed in order to write to the BIOS chip. When you decide that you want to upgrade the BIOS, and you have downloaded the correct BIOS update, you remove the jumper and then run the BIOS update. After updating the BIOS you place the jumper back.

Now that you have an understanding of what the BIOS is used for and how to update the BIOS, it's time to move on to CMOS and the different computer settings that can be controlled through CMOS.

Understanding CMOS

Complementary Metal-Oxide Semiconductor (CMOS) contains the computer's inventory list and advanced setup options. It can be considered an inventory list because it contains a record of all the devices connected to the system, such as the floppy drive, the hard drive, memory, and so on. Not only does CMOS list the devices, but it also dictates their capacity — for example, whether the system has a 2GB or a 6GB hard drive.

During the boot process, the system compares its inventory list to what it detects during boot-up. If there are any discrepancies, the system typically gives you an error and takes you into CMOS right away, asking you to save the new change. For example, assume that your system has 512MB of RAM and you add another 512MB of RAM to it. When you power the system up, the system compares what it had in inventory (512MB) the last time it booted with what it sees now (1024MB). Because there is a difference, CMOS reports

a memory size error and gives you the opportunity to save the new changes to CMOS. CMOS doesn't know that the difference is for the better; if it sees a difference, it reports an error. Because CMOS has detected the new memory, all you have to do is save the changes to the inventory list and reboot.



Looking at the previous example, if you choose not to save the settings you will get the memory size the next time you start the computer. In order to prevent the error from appearing again you will save the changes to CMOS — this way the values that are detected are equal to what is stored in CMOS.

Before discussing the different CMOS settings, it's important to know how to enter the CMOS setup program. Figuring this out is difficult on some systems, but extremely easy on others. Some systems display a message on boot-up that indicates what keystroke (often Delete or F1) to use to enter the CMOS setup program. IBM systems typically use the F1 key (refer to Figure 4-2); Dell systems today typically use F2 or the delete key.

Compaq systems typically use the F10 keystroke. In many Compaq systems, pressing F10 when you see a squared cursor in the top-right corner of the computer screen when the system starts up lets you enter the CMOS program.


In some older systems, you enter CMOS by holding down multiple keys at the same time during boot-up — Ctrl+Alt+S or Ctrl+Alt+Insert, for example. Entering into these systems is a little trickier, so reviewing your documentation is helpful. (Note that it has been a long time since I have seen a system that uses these three-key combinations to enter the CMOS setup program.) Table 4-1 summarizes popular keystrokes used to enter CMOS based off the manufacturer of the BIOS. (Remember that the CMOS setup program is stored in BIOS.)

Table 4-1 **How to Enter CMOS Setup**

<i>BIOS Manufacturer</i>	<i>Keystroke</i>
AMI	Del
Award	Del or Ctrl-Alt-Esc
Compaq	F10
Dell	F2 or Del
HP	F1
IBM	F1
Phoenix	F1 or F2

CMOS information is held in CMOS RAM, which is volatile memory that is able to maintain its data during shutdowns or power loss by using a small battery located on the motherboard. Thus, if the battery on the motherboard loses power, the CMOS data is also lost. Figure 4-3 shows a CMOS battery on a system board.

Figure 4-3: The CMOS battery on the system board is responsible for powering CMOS RAM.



The CMOS configuration is stored in CMOS RAM. Because RAM loses its contents when the system is powered off, the motherboard has a small battery that maintains enough of a charge that CMOS RAM can maintain its data — thus allowing the system to retain the CMOS information between reboots.

Viewing Basic CMOS Settings

In this section, you find out about common CMOS setup options and their purposes. These CMOS settings are consistent with most systems and have become fairly standard settings to view or change when troubleshooting a system.

After you have entered the CMOS setup program, you see an introductory screen, such as the one shown in Figure 4-4, that displays options that can be changed to control the system configuration.

Hard drive

Within the CMOS setup program you will be able to find out the size of the hard drive. CMOS displays the size of the hard drive based on what was detected during startup. The hard drive size is displayed in CMOS, as shown in Figure 4-5.

On some older systems you could change the CMOS from autodetect to manual so that you can specify the *dimensions* of the drive — overriding what was detected. Specifying the dimensions of the drive involves looking

on the back of the drive to find out how many sectors, cylinders, and heads the drive contains — which dictates the size of the drive. When you know the dimensions, you can input these dimensions into CMOS (after switching from autodetect to manual). CMOS calculates the size of the drive based on the dimensions you input.

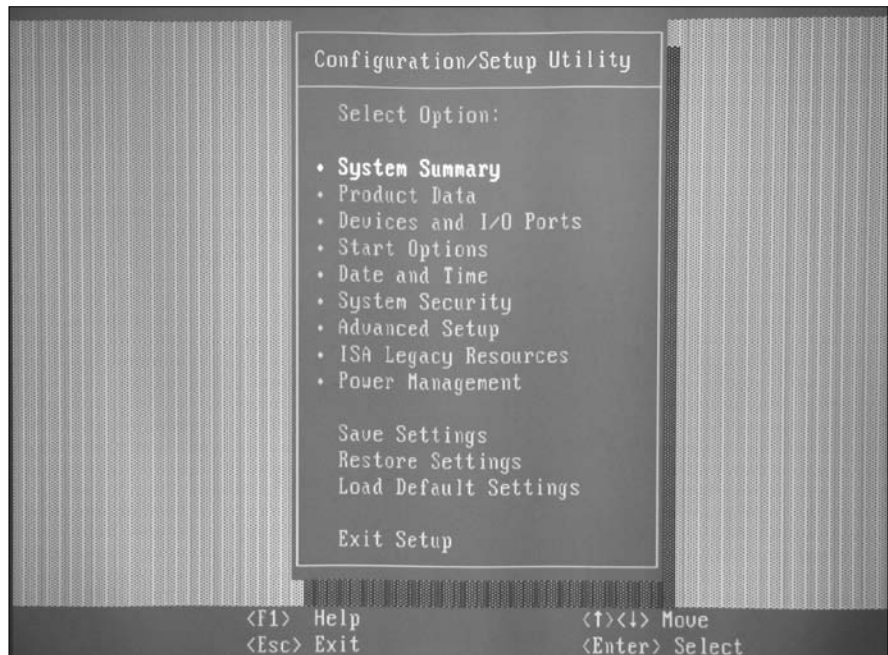


Figure 4-4:
A typical
CMOS
screen.




Figure 4-5:
Viewing the
hard drive
size in
CMOS.

On older systems, you specify the hard drive size by specifying a *drive type*. The drive type is a number representing a drive of a specific size. For example, a Type 2 drive may be 1.2GB in size, while a Type 10 drive may be 1.5GB in size (the actual values of hard drive types vary from system to system and are usually displayed along with the type). There is also a custom type on

older systems, usually Type 47, which enables you to specify the dimensions of the drive (sectors, cylinders, heads); CMOS can then calculate the size in megabytes for you after you input the dimensions.

Floppy disk drive

One option in CMOS lets you enable or disable the floppy disk drive. When the disk drive is enabled, its size is also specified and can usually be changed. If the disk drive is enabled but the wrong size disk drive is specified, you may not be able to access a floppy disk that has been placed in the drive. CMOS allows you to disable devices such as disk drives and USB ports so that companies can keep employees from copying proprietary data. Figure 4-6 shows the menu setting that allows you to change the size of the disk drive and even disable the disk drive.

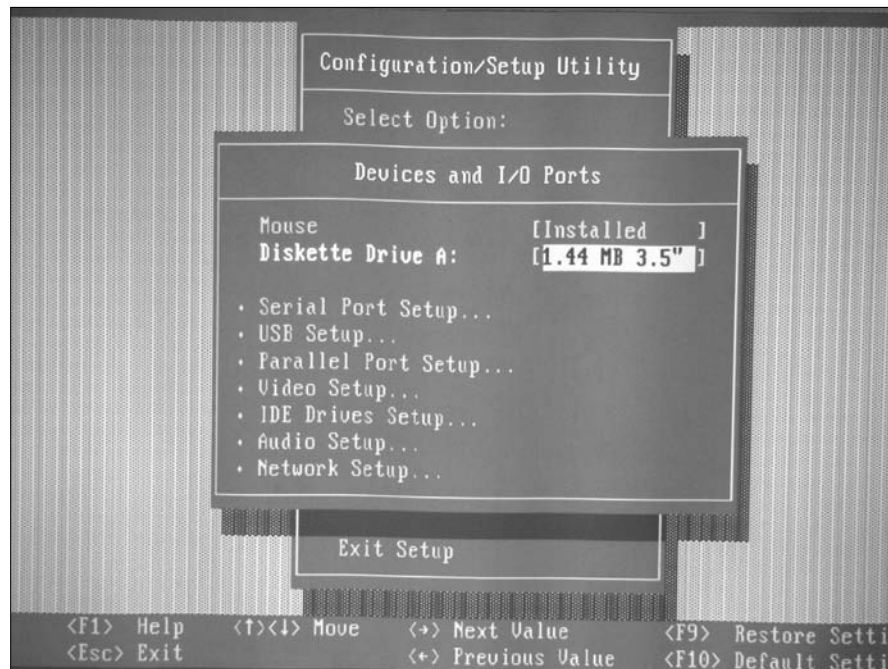


Figure 4-6: Changing the size of the floppy disk drive in CMOS.

Memory

CMOS indicates the total amount of memory installed on the system. Typically, this entry is not modified unless you add or remove RAM and the change is detected on boot-up. If you do add or remove RAM from the system, the system detects the change and modifies CMOS for you. All you need to do is save the modification once in the CMOS setup program. Some

systems today inform you that the system has saved the change for you and don't require you to enter the CMOS setup program. Figure 4-7 shows the CMOS System Summary screen, which shows that there are 640MB of system memory installed on the system.



Figure 4-7: Viewing the amount of memory installed on the system through CMOS.

The CMOS in some systems not only tells you how much memory you have installed but also informs you of the *type* of memory that is installed. These systems usually allow you to change the type of memory if the system supports different types of memory (for example, FPM, EDO, SDRAM, DDR; see Book II, Chapter 3 for more on memory types). For example, I have an old Pentium that supports either FPM or EDO memory, and CMOS displays the memory type being used.

Parallel ports

CMOS offers the opportunity to configure your parallel port. This configuration allows you to either disable the port or change the mode the port runs in.

You might think that a parallel port doesn't require a lot of configuration, but think about this: Have you ever had problems with a scanner that plugs into your parallel port? Or more specifically, have you ever had problems with your parallel-port scanner plugging into your computer and the printer then

connecting to your scanner? The problems may derive from the parallel port mode being misconfigured on the system. The *parallel port mode* dictates the capabilities of the parallel port. Table 4-2 lists the three parallel port modes that can be configured for your system.

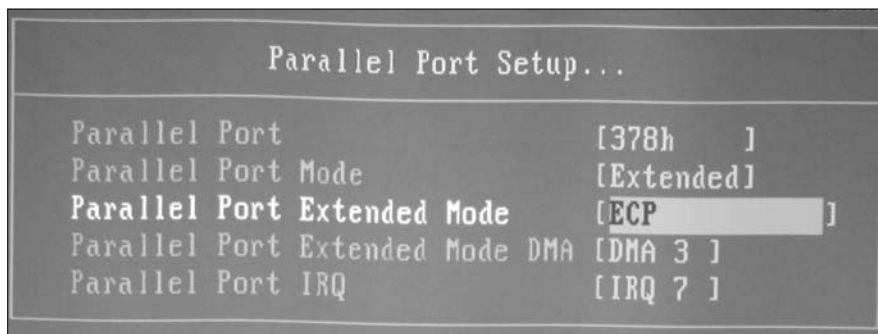
<i>Mode</i>	<i>Data Transfer Rate</i>	<i>Description</i>
Standard Parallel Port (SPP)	150 KBps	Supports communication in only one direction — from the computer to the device.
Enhanced Parallel Port (EPP)	2 MBps	A bidirectional port mode that enables communication in either direction — from the computer to the device and the device to the computer. This mode supports daisy-chaining (where one device is connected to another device that is then connected to the computer).
Extended Capabilities Port (ECP)	Over 2 MBps	Supports bidirectional devices.

If you're having trouble daisy-chaining the scanner and printer off the parallel port on the computer, check to make sure that the proper port mode is selected in the CMOS setup program (shown in Figure 4-8) — in this example, you want to make sure that EPP is selected because it supports daisy-chaining.



Daisy-chaining is the feature of connecting one device off another, such as connecting the printer to the scanner and the scanner to the LPT port.

Figure 4-8: Configuring the parallel port mode in CMOS.



The resources of the parallel port may also need to be configured if there are conflicts with another device. Notice that the default IRQ for LPT1 is 7, and the default I/O address is 378-37F. You will learn about IRQs and I/O addresses in Book III, Chapter 4 — for now, just make a mental note that you can configure the IRQ and I/O address of the parallel port in CMOS.

Serial ports

A typical system includes two serial ports, known as COM1 and COM2, and CMOS should have an entry for each of the serial ports — these entries enable you to change the resources, such as the IRQ (Interrupt ReQuest) and I/O address used by each of the serial ports, as shown in Figure 4-9. The CMOS shown in Figure 4-9 identifies the two serial ports as serial port A and serial port B. For more information on IRQs and I/O addresses, see Book III, Chapter 4.

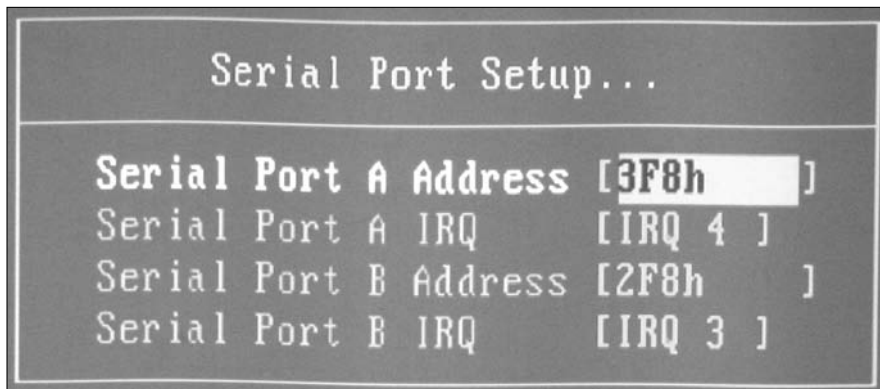


Figure 4-9: Configuring the serial ports in CMOS.



Be sure to remember that the default IRQ for COM1 is 4, and the default IRQ for COM2 is 3. The default I/O address for COM1 is 3F8-3FF, and the default I/O address for COM2 is 2F8-2FF.

In the CMOS setup program, you will not only find an option to change the resources for the serial ports, but you will also be able to disable the serial ports. Disabling the serial ports involves entering your CMOS program and switching a serial port to the *disabled* setting.

Date and time

The date and time each have an entry in CMOS as well, which is where the operating system gets its date and time information. If you set the date and time in CMOS it will set the date and time for the operating system. From a troubleshooting point of view, you know that your CMOS battery is dying when, during startup, the system asks you for the date and time. During startup after a battery failure, all of the other settings are detected again — which is why you don't specify the hard disk or the floppy disk. However, the date and time must be specified again.

Boot sequence

Take special note of the CMOS entry for boot sequence; it determines what devices the system will try to boot off of and in what order it tries each device. For example, older systems are typically set up to boot off a floppy disk first, and if a bootable floppy disk isn't present, the system then boots off the hard drive.

Newer systems let you boot off a CD-ROM device or even off the network first and then, if no bootable device is found in either of those spots, proceed to boot off the hard drive. Booting off a CD-ROM makes installing an operating system extremely easy because most operating systems today (like Windows or Linux) support booting off the installation CD to install the operating system.

Most systems today are configured to try to boot from a floppy disk, then a CD-ROM, then a hard disk, and finally a network. You determine the order to suit your needs, although typically, the floppy or CD-ROM drive is checked before the hard drive. If you like, you can disable devices, such as a CD-ROM drive — from being bootable. This is an important point because system security can be bypassed if a hacker can boot off his own CD containing his own copy of an operating system. For security reasons, your company may consider disabling booting from floppy disk or CD-ROM. Figure 4-10 shows the CMOS menu where you can configure the startup order for your system's devices.

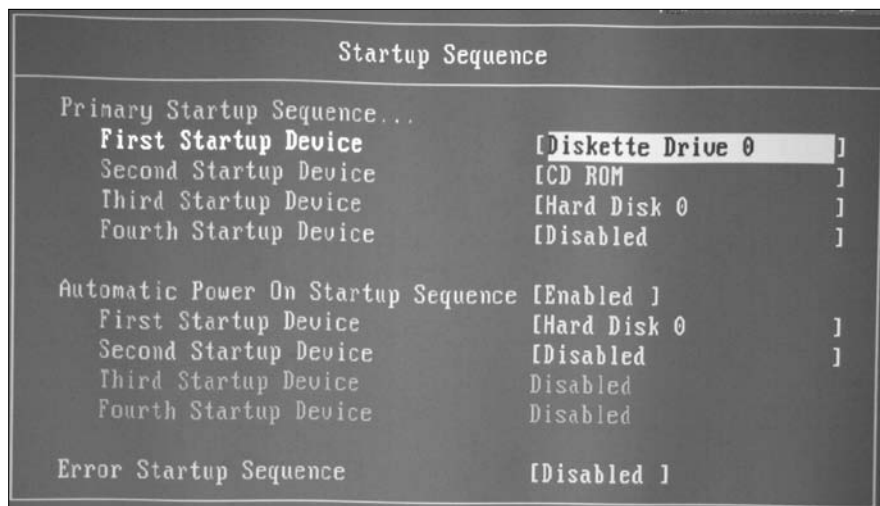


Figure 4-10: Configuring the startup sequence in CMOS.

Passwords

In general, two types of passwords can be set in CMOS: the power-on password and the administrator password. The steps to set the password are different per manufacturer but you will usually find the password option under a security menu in the CMOS setup program.

The *power-on password* is required in order to power the computer on. It is part of the power-on process and occurs before the operating system is loaded. Many people like to call it a *hardware password* because the operating system (the software) won't have a chance to load unless the correct password is typed in. The implementation of power-on passwords may be especially useful in environments in which security is a significant issue.

You set the *administrator password* in the CMOS setup program. This password is required for anyone wishing to enter the CMOS setup program and make changes to the system configuration. It prevents unauthorized users from entering the CMOS setup program and changing the values that reside there.


When people get comfortable with computers, they start to explore the computer's options. Companies often end up spending time and money fixing problems that arise from the exploration of these options. It may be useful to set an administrator password — that is, a password someone has to supply before entering CMOS to make changes. While the administrator password enables you to secure the workstation at the administrative level, the power-on password allows you to secure the system at the user level — controlling who can use the system.

A question I often receive is, "What happens if you forget the CMOS password after it is set?" The answer is simple. First, check to see if the motherboard has a jumper that can be removed to make CMOS forget the passwords (see Figure 4-11). Once you remove the jumper to clear the password you put the jumper back on and power on the computer. Some systems will have you place the jumper over specific pins in order to clear the password — check the documentation to find out how to clear the password for your particular system.

If no such feature exists on your motherboard, you could remove the CMOS battery from the motherboard. Remember that CMOS RAM retains its information because there is a battery supplying power to the memory where CMOS data is stored. If you remove the battery, CMOS is erased, including the passwords (it's better than throwing the system away!).

The only problem with removing the battery is that all the CMOS information is lost — including your hard drive, memory, and other device settings. The good news is that most of the information should be detected again on startup, such as the amount of memory installed or the size of the hard drive.

Figure 4-11: A jumper that can be removed to erase the CMOS password.



Earlier in this chapter, I mention that if your system asks you for the date and time when it boots up, this is usually an indication that your CMOS battery is dead or dying. As noted here, the date and time aren't the only things lost to a dying battery. If security is an issue, and if your system keeps asking you for the date and time at boot-up, replace the battery *immediately* and then reset the administrator password to keep others, malicious or benign, from altering the CMOS setup.

A third solution to a forgotten password — and the one I like — is to get a CMOS utility, such as the CMOS Save & Restore utility. CMOS Save & Restore backs up the CMOS information and then allows you to restore it when an emergency arises. This solution involves backing up CMOS *before* a password is set, so that CMOS can be restored to a state without a password. These utilities are popular ones that can be found on the Internet, and they back up the CMOS information to a text file. An example of a CMOS save and restore program can be found at <http://mindprod.com/products1.html#CMOSSV>.

Plug and Play BIOS

Most systems today have a *Plug and Play BIOS (PnP BIOS)*. The term *Plug and Play* refers to the idea that you can connect a new device to the system and the system automatically detects and configures the device to work with your system. To have a Plug and Play system, three conditions are necessary:

- ◆ You need a Plug and Play device.
- ◆ You need a Plug and Play operating system.
- ◆ You need a Plug and Play BIOS.

If you're missing any one of these conditions, the operating system will not be able to leverage Plug and Play and assign resources to a device on startup. On some older systems you may see an entry in CMOS stating that it is a Plug and Play BIOS. All newer systems are Plug and Play BIOSes.

Viewing Advanced CMOS Settings

Many of the newer systems maintain the basic CMOS parameters mentioned in the previous section but are also supplemented by different advanced setup settings. This section examines the purposes and characteristics of some of these advanced settings.

Globally Unique Identifier (GUID)

The Globally Unique Identifier (GUID) is a 128-bit number randomly generated for the system when it was built and is stored in CMOS. The GUID uniquely identifies the system from any other. It enables the identification of individual computers and ensures that this identification method is 100-percent unique.

Many features in today's computers make use of the GUID. For example, one of the Windows installation tools (Remote Installation Services, or RIS) allows the administrator to go into a client computer's CMOS, make a note of the GUID, create a computer account within Windows Active Directory database for the workstation, and associate the GUID with the account. If the client computer is booting off the network when it starts, it will contact the RIS server and start installing a Windows client operating system. The client computer will use the computer name of the computer account found in Windows Active Directory that its GUID is associated with.



Some systems have adopted the term *UUID* (Universally Unique Identifier) instead of GUID. Remember, you may see a UUID in CMOS instead of a GUID, as shown in Figure 4-12, but they're the same thing.

BIOS date and revision number

The date of your BIOS should be displayed somewhere in CMOS, usually under Summary Information. If you don't have a BIOS date, you may have a revision number or level. These entries in CMOS are important because you may be required to update your BIOS someday by going to the manufacturer's Web site and downloading the update. The first thing you will notice when doing this is that the manufacturer has built many different versions of BIOS for its systems. You need to make sure you get the proper update (the one for the date of your BIOS) by watching the revision number of your BIOS. For example, the site may tell you to download BIOS update "1234" if you have Revision number R5.145. The big question is how do you know what your revision number is? Check CMOS! Refer again to Figure 4-12. You can see the BIOS version, or revision number, identified by the label "Flash EEPROM Revision Level."

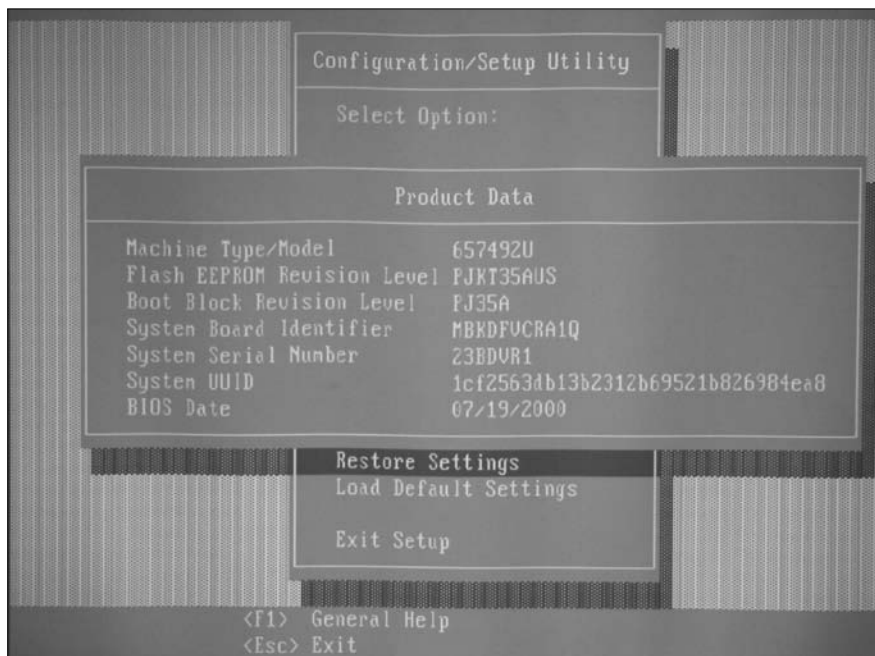


Figure 4-12: Viewing the system UUID and BIOS revision number through CMOS.



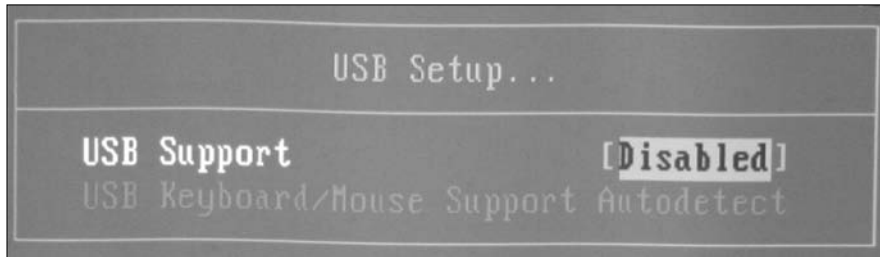
To ensure that you have downloaded the correct BIOS update for the system it is always a good idea to make a note of the revision number and the date of the existing BIOS. Both pieces of information are typically used to determine which BIOS you need to download from the manufacturer's Web site.

Universal Serial Bus (USB)

Universal Serial Bus (USB) devices have gained much recognition over the last decade. USB devices are high-speed serial devices that use a single connector style and can be chained together with a *USB hub device*. A USB hub device connects all USB devices together at a central point. The USB hub may be its own unique device, or it may be just another device in the USB chain that has the capability of connecting other USB devices to it. For example, a USB monitor may have a USB port to allow a mouse to connect to it. Some popular USB-type devices include digital cameras, scanners, mice, and keyboards. You can even find USB network adapters.

If you have any problems getting a USB device connected to your computer, make sure the USB port has been enabled in CMOS. Also, remember that you can disable the USB ports in CMOS (as shown in Figure 4-13) to help secure your environment. You may wish to do this to prevent a user on the network from using the flash drive to take proprietary corporate data home.

Figure 4-13: Disabling the USB ports on a system prevents the use of flash drives or other USB-type devices.



Built-in network adapter

Most systems today come with built-in network cards that allow the system to connect to a network or the Internet. Because the card is built-in, you don't have to buy a network card for the system. However, built-in network cards (or any built-in device) sometimes become faulty, but your system doesn't know that and always tries to use this built-in device anyway. Note that you can usually enable or disable the built-in network card through the CMOS setup program.

When your system includes a built-in network adapter, you usually have the option of enabling or disabling the capability to boot off the network as well. Unless this option is enabled, you will not be able to boot off the network — even if you have specified a network boot in your startup order.



You need to verify three options in CMOS in order to boot off the network: verify that the network card is enabled, verify that your system has enabled booting off the network adapter, and then ensure that the network device is located in the startup order of devices. To boot off the network, you probably want the network as the first boot device.

Virus protection

Some BIOS systems have built-in virus protection — which is, for the most part, a good thing. Some viruses attack the system by altering the *Master Boot Record (MBR)*. The MBR is the code that initiates the startup of the system and is located at the beginning of the hard drive. For more information on MBR and hard drives, check out the next chapter in this minibook. The virus protection built into the BIOS watches out solely for changes to the MBR and puts a stop to it!

Unfortunately, when you install a new operating system, the installation modifies the MBR with its boot program files. As a result, the built-in virus detection in CMOS will see that “something” is modifying the MBR and will assume it’s a virus and stop the new operating system from loading. In this case, you need to go into CMOS, disable the virus protection, and then restart the installation. Figure 4-14 shows the startup options that indicate that virus protection has been disabled.

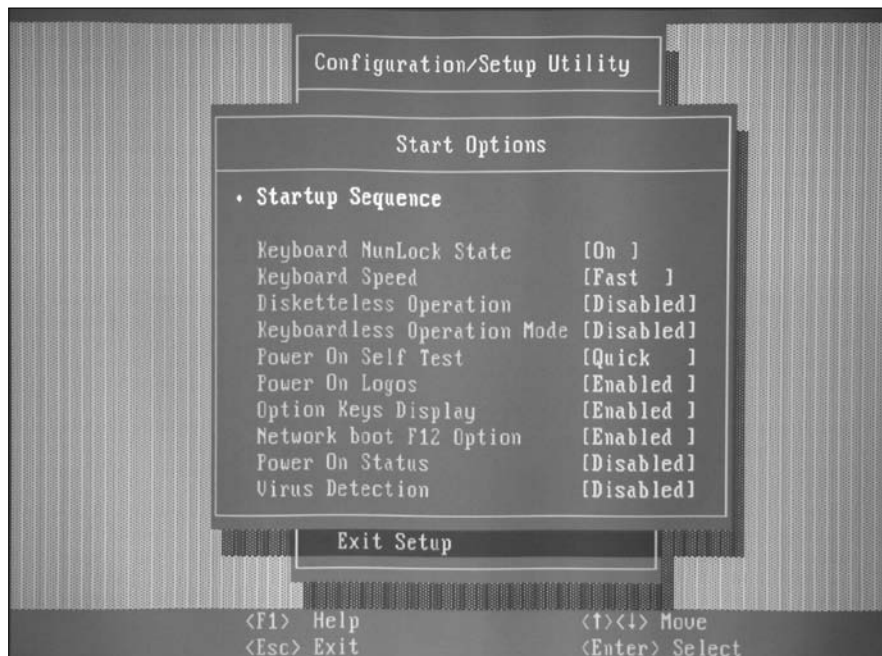


Figure 4-14: Viewing the virus protection setting in CMOS.


On-board cache

In Book II, Chapter 3, I discuss the benefits of cache and the types of cache memory, L1 and L2. In CMOS, you typically find an entry indicating how much cache memory exists on the system, and you can configure CMOS to disable this built-in cache memory.

Sometimes cache memory goes bad and causes boot-up problems. If having cache memory enabled presents any compatibility problems with your system, you can try disabling the cache memory as a troubleshooting technique. Figure 4-15 shows the size of the cache memory and the option to

enable or disable the cache memory. If the problem doesn't go away, then there was no problem with the cache memory to begin with; enable the cache memory once again.

Figure 4-15:
You may disable built-in cache memory in CMOS to troubleshoot problems with corrupt cache memory.



Reserve resources

Many systems let you reserve resources that have been assigned to *legacy devices* in CMOS. Legacy devices are non-Plug and Play devices, such as ISA (Industry Standard Architecture) cards, that have been added to the system. The ISA non-Plug and Play devices are hard coded with a particular I/O address and IRQ, so you may need to configure CMOS to reserve the IRQ so that it is not given to a Plug and Play device on startup.

Removing the IRQ will ensure that a Plug and Play system will not assign the resources that are hard-coded into older devices to Plug and Play devices, which would create a conflict. Most systems today are Plug and Play, so, unless you're supporting older hardware, you will most likely not hit this issue. Figure 4-16 shows how to reserve a resource in CMOS by setting the resource to "ISA" — meaning it is being used by an ISA device, so the system is not to give this resource out to another device.



Now it is time to get some hands-on by performing Lab 4-1 and Lab 4-2! Lab 4-1 and Lab 4-2 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Interrupt Resources	
0:	Timer
1:	Keyboard
2:	Interrupt Controller
3:	Serial Port B
4:	Serial Port A
5:	[ISA Resource]
6:	Diskette
7:	Parallel Port
8:	Real Time Clock
9:	ACPI
10:	[Available]
11:	[Available]
12:	Mouse
13:	Coprocessor
14:	IDE Drives
15:	IDE Drives

Figure 4-16: Reserving a resource in CMOS so that the system does not give the resource out.

Getting an A+

In this chapter, I introduce you to the common settings found in CMOS setup programs today. You discover the basic and advanced parameters in CMOS. The following points are covered:

- ◆ The BIOS is low-level code used by the system to communicate with the system hardware. The BIOS code is stored on an EEPROM chip located on the motherboard of the system.
- ◆ The BIOS is typically updated by downloading a BIOS update program and “flashing” the BIOS.
- ◆ You enter a system’s CMOS setup program by using a keystroke such as F1 (IBM), F2 (Dell), or F10 (Compaq) when the system starts up.

- ◆ You can change the boot order of the different bootable devices.
- ◆ There are three parallel port modes — SPP, EPP, ECP.
- ◆ You can change the hard drive from autodetect to manual so that you can input the drive dimensions.
- ◆ You may update your BIOS by downloading the BIOS update for your system from the BIOS manufacturer's Web site.

Prep Test

- 1** You have a USB camera from which you are trying to copy some of the pictures to the computer's hard drive. The computer does not seem to recognize the device. What should you do?
 - A Check to see that the boot order is correct.
 - B Make sure that the USB ports are enabled.
 - C Make sure that the serial port is enabled.
 - D Make sure that the operating system has not assigned the USB device resources to some other device.

- 2** The computer constantly prompts for the date and time. What does this indicate?
 - A The date and time are wrong.
 - B The time has fallen back by one hour.
 - C The battery on the motherboard is losing its charge.
 - D The time has increased by one hour.

- 3** On a Plug and Play system in which you have some non-Plug and Play devices, what should you do so that there are no conflicts with the legacy hardware resources and the resources assigned by the operating system?
 - A Replace the BIOS with a non-Plug and Play BIOS.
 - B Enable the Legacy Compatible option in CMOS.
 - C Replace all Plug and Play devices with non-Plug and Play devices.
 - D Go into the reserve resources area of the CMOS setup program and reserve the resources used by legacy devices.

- 4** Your manager is worried that someone will be able to start up any of the computers and view confidential information. What could you do to ensure that anyone starting the computer is supposed to be using that system?
 - A Lock the computer case.
 - B Set file permissions on the files that are confidential.
 - C Set a power-on password in CMOS.
 - D Use Microsoft Encrypting File System to protect the files from unauthorized access.

- 5** You are installing Windows on your computer and you know that the Windows CD is a bootable CD. You have tried a number of times to boot off the CD, but you are unable to. What CMOS option should you look for?
- A You must delete the existing partitions so that you can boot off the CD-ROM.
 - B Ensure that the CD-ROM device has been set up as the first device in the startup (boot-up) order.
 - C Disable the hard drive in CMOS.
 - D Disable the CD-ROM in CMOS.
- 6** You are worried that some of your advanced computer users on the network will start changing the settings within CMOS. What is the best thing you can do to protect these settings?
- A Set a power-on password.
 - B Set an administrator password.
 - C Set a Windows log-on password.
 - D Ask the advanced users not to change any of the CMOS settings.
- 7** When looking inside the My Computer icon, you notice that no floppy drive is listed. What should you check in CMOS?
- A Make sure that the floppy drive is listed as a bootable device.
 - B Make sure that the hard drive is listed as a bootable device.
 - C Make sure the floppy drive is configured correctly and enabled in CMOS.
 - D Make sure the hard drive is configured correctly in CMOS.
- 8** Your built-in network adapter doesn't seem to be connecting you to the network. What is one of the first things you want to check for in CMOS?
- A That the built-in network adapter is enabled in CMOS.
 - B That the network adapter driver is loaded.
 - C That the resources are not conflicting.
 - D That the proper protocol is installed.
- 9** You are trying to use Windows 2000 Server's Remote Installation Services, but the computer doesn't seem to boot off the network adapter. You have verified that the network adapter has been enabled in CMOS. What else should you check for within CMOS?
- A That the network card has been disabled.
 - B That the network adapter has been disabled.
 - C That the CD-ROM has been configured as a bootable device.
 - D That the network adapter is set up as a bootable device.

- 10** You are experiencing a lot of problems running your customized software on some of the newer computers. What might you try disabling in CMOS to clear up the compatibility issue between the software and the system?
- A RAM
 - B Cache memory
 - C Hard drive
 - D Floppy drive
- 11** You are having trouble installing a new operating system on your computer because the installation fails with each attempt. Which of the following CMOS settings would you disable to help the system make it through the installation?
- A Network adapter
 - B Bootable CD-ROM
 - C Antivirus
 - D Bootable network adapter
- 12** Which of the following is a typical method for updating your system BIOS?
- A Replacing the old BIOS chip with a new BIOS chip
 - B Running a BIOS update program from a diskette that flashes the BIOS with a new version of the BIOS
 - C Replacing the old motherboard with a new one
 - D You cannot update the BIOS
- 13** Which of the following best describes the difference between BIOS and CMOS?
- A BIOS contains the configuration information for the system, while CMOS is the low-level code that allows the devices to communicate.
 - B BIOS is stored in RAM, while CMOS is stored in ROM.
 - C BIOS is stored in RAM, while CMOS is stored in cache memory.
 - D CMOS contains the configuration information for the system, while BIOS is the low-level code that allows the devices to communicate.
- 14** Which two of the following are typical keystrokes used to enter the CMOS setup program?
- A F1
 - B Ctrl+Alt+Delete
 - C Alt+F10
 - D Delete

Telling Your BIOS from Your CMOS

- 15** You are deploying Windows 2000 Professional computers by taking advantage of Remote Installation Services. What CMOS setting do you need to look up on the computer you are deploying the new operating system to?
- A GUID
 - B Network adapter address
 - C Serial number
 - D ISBN
- 16** Where is the CMOS configuration information stored?
- A RAM
 - B ROM
 - C Cache memory
 - D Floppy disk

Answers

- 1 B.** In today's systems, many of the built-in devices can be enabled or disabled in the CMOS setup program. When experiencing problems with a built-in device, the first thing you should check is whether the device is enabled. See *"Built-in network adapter."*
- 2 C.** The CMOS configuration information is stored in CMOS RAM. This special RAM chip maintains its information by using a small battery on the system board. If the battery loses its charge, then the CMOS RAM is flushed out, meaning that it will prompt you for the information such as the date and time as the system boots up. *Review "Date and time."*
- 3 D.** The newer CMOS setup programs allow you to reserve resources so that the OS will not give those resources away. If you reserve the resources, the system will not give those resources out. This way, when the legacy device initializes, the resource it tries to use is free. *Check out "Reserve resources."*
- 4 C.** Choice C is the best answer to ensure that only authorized individuals are powering on the system. In order to use the system, someone will need to type the CMOS password. This password is not dependent on any particular operating system because you are setting a hardware-level password. *Peruse "Passwords."*
- 5 B.** Many computers are set up to boot off the floppy drive, then the hard drive, and finally the CD-ROM, which means that if the system can boot off the floppy or hard drive, the opportunity will never arise to boot off the CD-ROM. Changing the startup order so that the CD-ROM is the first bootable device means that you can put the Windows CD-ROM in and setup will be invoked from it. Choice A would work, assuming that the CD-ROM device is listed anywhere in the bootable device order. Unfortunately, you have lost the contents of the hard drive because partitions have been wiped out, so Choice A is not the best answer. *Take a look at "Boot sequence."*
- 6 B.** An administrator password is a password that must be provided in order to enter the CMOS setup program and change the settings. Choice A is not the right answer because you may not want to set up a password for when the computer turns on but only if a user tries to enter the CMOS setup program. *Peek at "Passwords."*
- 7 C.** The floppy drive is a device listed in CMOS. Not only do you want to make sure that the floppy drive is physically connected correctly, but you also want to make sure that CMOS sees the device. Check to make sure that the floppy drive is not disabled in CMOS. *Look over "Floppy disk drive."*
- 8 A.** One of the first things you want to do with a built-in device is ensure that the device has not been disabled in CMOS. If it has not been disabled, you should look at the other choices for the solution. *Study "Built-in network adapter."*

Telling Your BIOS from Your CMOS

- 9** D. When taking advantage of Remote Installation Services, you need to ensure that the computer not only has the network adapter enabled, but also that the network adapter is signified as a bootable device, and maybe even the first bootable device. In this question, you have already verified that the network adapter has been enabled, so Choice D is the only possible answer. *Refer to "Built-in network adapter."*
- 10** B. Many systems today have built-in cache memory. Although it isn't that common, you may sometimes have problems with particular software not "liking the idea" of using information from cache instead of RAM. To test this, you may temporarily disable the cache memory. *Examine "On-board cache."*
- 11** C. Viruses commonly attack the MBR of the hard drive, so virus software constantly watches out for applications that try to make a change to this area of the disk. Some systems today have built-in virus detection, and because the installation of a new operating system causes a change to the MBR, there may be incompatibilities between the operating system installation and the virus protection built into the system. Disabling this virus protection in the CMOS will allow the installation to finish. *See "Virus protection."*
- 12** B. To update your BIOS today, you will either get the update from the manufacturer on a floppy disk or download the update from the Internet onto a floppy. You will then boot off this diskette, and the update process will start. In the past, you would have to replace the entire ROM chip. *Review "Performing the BIOS upgrade."*
- 13** D. CMOS is the "inventory list" of devices and their configuration, whereas the BIOS is the set of low-level instructions that tells these devices how to communicate. Choice B and Choice C are wrong because BIOS code is stored in ROM, not RAM. *Check out "The BIOS and Its Purpose."*
- 14** A. and D. Typical keystrokes used to enter CMOS while booting are F1 and Delete. *Peruse "Understanding CMOS."*
- 15** A. When taking advantage of Remote Installation Services for Windows 2000, you will build a computer account in Active Directory and associate a GUID with the computer account. The GUID is located in CMOS and is a unique number associated with that system. *Take a look at "Globally Unique Identifier (GUID)."*
- 16** A. CMOS configuration information is stored in RAM, which maintains its information while the power is turned off by using a small battery located on the motherboard. *Study "Understanding CMOS."*

Chapter 5: Working with Storage

Exam Objectives

- ✓ Understanding hard drive terminology
- ✓ Discovering IDE devices
- ✓ Learning how SCSI works
- ✓ The Serial ATA hype
- ✓ Using removable storage
- ✓ Working with file systems and partitions

One of the primary responsibilities of a PC technician is to manage hard disks. This responsibility involves more than just partitioning and formatting disks; it entails installing disks into the computer system and configuring the system to recognize the newly added disks.

In this chapter, you discover the skills that you need for installing, configuring, and troubleshooting hard disks. We introduce you to a new world of terminology related to hard disk installation and configuration, along with other types of storage like CD-ROMs and DVDs.

In this chapter, We also make sure that you understand the difference between IDE, SCSI, and Serial ATA devices. You also see the steps required to install and configure IDE master/slave devices in a chain and the steps required to install a SCSI bus.

Understanding Hard Drive Terminology

If you took a martial arts class, the instructor would show you how to do some basic punches and kicks before letting you spar or compete. The instructor would be well aware that jumping into combat without the basics could do you more damage than good. PC repair works the same way: Before discussing the installation and configuration of IDE and SCSI devices, it is important to cover some basic terminology surrounding hard drives.

Disk geometry

The following sections introduce you to *disk geometry* — essentially the physical components of a drive that make up your data storage solution.

You also find out some general terminology about hard drives and hard drive storage in this section.

Platters

A *platter* is a physical object (actually, a plate) inside the hard disk that is responsible for storing the data. A platter is similar to a record on an old record player — the main difference being that a hard disk has many platters, while a record player holds only one record at a time. Figure 5-1 shows the platters inside a hard disk.




Figure 5-1:
The
internals of
a hard disk.

The platters are very much like records on a record player in the sense that they spin on a spindle that runs through the center of all the platters. Each platter has two sides for storing information, and each side of the platter has a unique ID. The ID for the first side of the first platter is 0, and each side increases by 1. For example, if there were two platters in the disk, the first platter would have Side 0 and Side 1, while the second platter would have Side 2 and Side 3. Figure 5-2 illustrates this concept.




Figure 5-2:
Sides of a
platter.

Because each side of the platter has a writing mechanism, many people use the terms “head” and “side” interchangeably. The head is more accurately called the *read/write head* because it moves over the disk surface and reads from or writes to the disk. Like a needle on a record player, the read/write head moves over the surface of the disk with the help of an arm, called the *actuator arm* or the *head positioning mechanism*. Figure 5-1, shown previously, illustrates the read/write head on an actuator arm.



Each platter surface on the disk has its own read/write head. When information is written to the disk, the read/write head moves to the same track on all platters in a single movement and then writes the data across the same track on all platters. The actuator arm has multiple read/write heads on it.

Tracks

Just as there are grooves, or tracks, on a music record, there are also *tracks* on each platter. These tracks are evenly spaced across the platter’s surface.

Sectors

The platter is divided into pie-shaped slices, called sectors. Now the confusing thing about sectors is that where a track intersects with a sector creates *sector blocks* — also known as sectors! Each sector (block) is 512 bytes in size, and is the actual storage area for data. Figure 5-3 illustrates the tracks and sectors on a disk platter. Each pie-sliced sector has an address; the first sector is known as sector 1, the second sector is known as sector 2, and so on. The reason why the term sector applies to two different areas of the disk is because if you were to look at a sector (the pie slice), like the pie slice labeled in Figure 5-3 as “sector,” it may be sector one. Now all of the sector blocks in that pie slice are all considered sector one as well, but they differ in the track that they reside on. So each sector block has an address that is made up of the platter side number, the sector, and track number. For example, data can be saved to side 1, sector 2, track 4 — which is the address of a 512-byte sector block. Note that the term sector block is a term I made up for this discussion; the term *sector* is also used to describe the 512-byte blocks.




Figure 5-3: Tracks and sectors on a disk platter.

Clusters

A group of sectors makes up a *cluster*, and a cluster is the allocation unit for a file — meaning where a file is saved. When a partition (a partition is a logical division of space on the disk; refer to the section titled “Managing Partitions and Volumes,” in this chapter) is formatted, the file system determines the cluster size based on the partition size. For example, a 2GB FAT partition uses a 32K cluster size. That same 2GB partition formatted as FAT32 uses only a 4K cluster size.

Having a partition use a 4K cluster size means that eight sectors make up a cluster. Keep in mind that after a file has been saved to the cluster, no other file can occupy that cluster. For example, if you have a 32K cluster size and you save a 3K file to the hard disk, the file is saved to an empty cluster — but only 3K of that cluster is used, and the remaining 29K is empty. The remaining 29K is now considered unusable space; no other file can be saved to that unused 29K.

Cylinders

All the platters in the hard disk contain the same number of tracks, but that number varies from one hard disk to another. These tracks are numbered from the outside in, starting with zero. For example, on a platter with 10 tracks, the track closest to the outer edge of the platter is Track 0, while the track closest to the center is Track 9.

A *cylinder* consists of the same track on both sides of all the platters. In other words, when you reference Track 0, you reference a particular track on a particular platter, but when you reference Cylinder 0, you reference Track 0 on all platters.

If you know the number of cylinders, heads, and sectors per track, you can calculate the size of a disk. For example, if a drive has 4,092 cylinders, 16 heads, and 63 sectors per track, the size of the disk would be 2,111,864,832 bytes (2.1GB). The formula to calculate the size of the disk is number of cylinders \times number of heads \times number of sectors per track \times 512 bytes per sector.

Read/write process

Platters are divided into 512-byte sectors. These sectors are the area on the platter that data is written to. The platters have a magnetic coating applied to them that is extremely sensitive to magnetism. While the platters spin, the read/write head moves from track to track until it reaches the desired track. Then it waits for the appropriate sector to move underneath it, at which time the read/write head is energized to apply a magnetic charge to the particles in the disk coating. This changes the particle binary state from zero to one, thus creating data. The same happens when the data needs to be read, the read/write head moves over the appropriate sector and reads the data that resides in the sector.

The read/write heads don't actually touch the surface of the disk platters; instead, they hover about ten micro-inches (or millionths of an inch) above it — that's not even enough space to place a hair between the read/write head and the platter's surface. This design helps improve disk performance because a read/write head that makes contact with the platter would cause friction, slowing down the rotation speed of the disk and creating extra heat.

Performance

Disk performance can be measured in terms of several important characteristics:

- ◆ **Seek time** is the time it takes to move the read/write heads to the desired track. Seek time is a calculated average because the time it takes to move to the desired track differs from one instance to another. For example, if the read/write heads are on Track 1, they will take longer to move to Track 12 than to Track 3. Seek time is measured in milliseconds, or one thousandth of a second. Figure 5-4 shows how seek time is measured.
- ◆ **Latency** is the time it takes for the appropriate sector to move under the read/write head. Latency is measured in milliseconds.
- ◆ **Access time** describes the overall speed of the disk. It is a combination of seek time and latency. The lower the access time, the better.
- ◆ **Spin speed** is the speed at which the platters spin, measured in rotations per minute, or rpm. The larger the rpm value, the faster the disk, which means less latency.




Figure 5-4:
Measuring
seek time.

Master Boot Record

The *Master Boot Record (MBR)* is the first sector on the first track of the first side of the first platter; it holds the operating system boot code that controls the loading of the operating system.

The MBR also holds drive characteristics, such as the partition table. During the boot process, the system has to find a primary partition that is active — it does this by looking at the partition table in the master boot record.

In general, if anything goes wrong with the MBR, you won't be able to boot the system. Because the boot record is always in the same location on every disk, it becomes very easy for a malicious hacker to write viruses that modify or corrupt the MBR. This is one reason you should always run virus-detection software. Refer to Book IX, Chapter 3, for more on virus detection and protection.

LBA and ECHS

For the A+ exam you need to be familiar with *Logical Block Addressing (LBA)* and *Extended Cylinder/Head/Sector (ECHS)* and what their purpose is. Essentially LBA and ECHS perform the same goal — they perform *sector translation*; sector translation is the hard drive controller lying to the BIOS about the drive geometry. LBA was developed by Western Digital while ECHS was Segate's solution to recognizing larger drives. Keep reading to learn why you need sector translation.

The reason you need sector translation is that the original BIOS code found on computers was limited to being able to see only 1024 cylinders, 16 heads, and 63 sectors — which is a total drive size of 504MB (1024 × 16 × 63 × 512).

The problem is if you bought a 2.1GB hard disk, your BIOS would not recognize it because the geometry of the 2.1GB drive is too high for the BIOS. In this example the geometry of the drive is 16,384 clusters, 4 heads, and 63 sectors. By looking at Table 5-1 you can see that the lowest value in each category is what will be recognized by the system.

Table 5-1	BIOS Limit Example			
	<i>Cylinders</i>	<i>Heads</i>	<i>Sectors</i>	<i>Total Size</i>
BIOS limit	1024	16	63	504MB
Hard drive (physical geometry)	16384	4	63	2.1GB
Detected	1024	4	63	132MB

As an example of why you take the lowest value in each category, if the hard disk only supports 4 heads then only 4 heads are detected. Although the BIOS supports a potential 16 heads, that doesn't mean they are actually there.

So the problem is that you have purchased a 2.1GB drive but the system is only recognizing 132MB! The solution to the problem is LBA or ECHS — again, both of these technologies offer the same solution. They were just built by different manufacturers.

An LBAenabled BIOS can recognize 1024 cylinders, 256 heads, and 63 sectors — essentially being able to support more heads on the drive. As a result, the drive lies to the BIOS by using a *translation factor* of usually 2, 4, 8, or 16. The physical dimensions of the drive are taken and manipulated by the translation factor to calculate the logical dimensions that are reported to the BIOS. In our example, 16,384 cylinders are too many cylinders so they are divided by translation factor of 16 to reach the LBA maximum number of cylinders supported. To make up for the loss in cylinders, the heads are then multiplied by 16, ensuring that the logical number of heads falls under the LBA limit of 254. Table 5-2 shows the solution — notice that the size of the drive (2.1GB) is what the LBAenabled BIOS will recognize.

Table 5-2 LBA Enabled BIOS Translation Example

	<i>Cylinders</i>	<i>Heads</i>	<i>Sectors</i>	<i>Total Size</i>
LBA enabled BIOS limit	1024	256	63	8.4GB
Hard drive (physical geometry)	16384	4	63	2.1GB
Translation Factor	Divide by 16	Multiply by 16		
LBA Translated Geometry	1024	64	63	2.1GB



In order to leverage larger size drives, your BIOS would have to support LBA or ECHS — which most BIOS do today.

Notice that an LBA enabled BIOS can only support an 8.4GB drive — and we are way past that drive size today. Today’s BIOS support the INT13 extensions, developed by Phoenix Technologies, which allow the systems to see drives past 137GB in size! The BIOS can recognize larger size drives because it simply identifies the drives by the number of sectors.

Discovering IDE Devices

In the following sections, you find out how to install and troubleshoot IDE devices. First, you get an overview of IDE devices and some of their features, and then you see a number of different configurations for installing IDE devices into a computer system.

IDE overview

The *hard drive controller* is responsible for converting signals made by the system CPU to signals that the hard disk can understand. These signals

include instructions on where to find data and how to get to that data. The hard disk performs its task, and any data that needs to be returned is sent to the controller from the hard disk. The controller then converts the signals from the hard disk into signals that the system can understand.


In the past, the controller was on an expansion card, which was connected to the drives via ribbon cables. The goal of IDE was to make the installation of hard disks easier by including the controller on the hard disk, which is where the name comes from: *Integrated Drive Electronics (IDE)*. So, today's drives have the controllers integrated into the drives themselves, meaning the drive is its own boss.

Originally, IDE was available only in the flavor of hard drives; originally you did not have any other type of IDE device such as CD-ROM or tape drives. IDE hard drives had a maximum capacity of about 528MB. Another important limitation with IDE is that only two devices could be connected in a chain. Back in the day when IDE was first used, SCSI was better in that respect: It supported eight devices in a chain (more on SCSI in the “Learning How SCSI Works” section of this chapter).

Original IDE devices have a transfer rate of about 10 Mbps and may have cache on the drive itself. The *cache memory* is a small amount of memory for storing data that is used frequently to increase drive performance.

Due to its limitations, IDE has been replaced by *Enhanced Integrated Drive Electronics (EIDE)*. EIDE devices have a transfer rate of about 16 MBps. Four devices are allowed in an EIDE chain, with a greater variety available. For example, you may now add CD-ROMs and even zip drives to the EIDE chain. Note that the capacity of the drives has been dramatically increased with EIDE — it now supports drives of over 200GB! Table 5-3 compares the features of IDE and EIDE.

Table 5-3		IDE versus EIDE
<i>Characteristic</i>	<i>IDE</i>	<i>EIDE</i>
Size	528MB	200GB and higher
Devices in Chain	2	4
Transfer Rate	10 MBps	16 MBps and above
Types of Devices	Hard drives	Hard drives, CD-ROMs



On the A+ exam, you may be asked the number of devices that IDE supports. IDE supports two devices, while EIDE supports four devices. Although there have been improvements in the IDE technology, Table 5-3 displays some of the original specifications.

IDE has been around for quite some time now, and as a result has gone through some changes. The following sections outline some of the technologies and terms that describe the different versions of IDE. Be sure to know these for the exam.

IDE/ATA standard

A number of hard drive standards have been developed over the last two decades — the first major standard being the IDE standard. The Integrated Drive Electronics (IDE) standard, which has been around since 1989, calls for an integrated controller on the drive to manage information entering and leaving the hard disk.

IDE drives attach to the motherboard by means of a 40-wire ribbon cable. The IDE standard also allows two drives to daisy-chain, creating a master/slave relationship between devices. The master drive is responsible for sending and receiving information in the chain.

The IDE standard is also known as the *Advanced Technology Attachment (ATA) standard*, which is sometimes known as the *ATA-1 standard*.

EIDE/ATA-2 standard

The *Enhanced Integrated Drive Electronics (EIDE)* standard followed shortly after the IDE standard. The EIDE standard allows four drives to be connected to a dual-channel controller. This is usually implemented as a motherboard with two connectors, also known as controllers, one primary and one secondary. You could then connect two drives off of each controller, making a master/slave chain for each controller. Figure 5-5 shows a primary controller and a secondary controller on the motherboard. EIDE also supports larger hard disks than the original IDE standard; the original size of an IDE drive was approximately 528MB. The EIDE standard is also known as *ATA-2 standard*.




Most techies in the computing industry interchange the terms IDE and EIDE. For example, if you have a newer motherboard that uses EIDE drives, most documentation and technicians simply use the term IDE — although the drives may be EIDE, or even an Ultra DMA drive. Also note that most techies and documentation use the term controller for the connectors on the motherboard that you connect the drives to — but really the controller is the circuitry on the hard drive that controls access to the drive. My point is get used to people using the terms IDE instead of EIDE and controller instead of connectors.

ATAPI specification

Originally, IDE devices were implemented as hard drives, but an additional ATA specification allows other types of devices to exist on an ATA (or IDE) chain. This specification is known as the *ATA Packet Interface (ATAPI)*, which

allows devices like CD-ROMs and tape drives to exist on an ATA chain. Other types of ATAPI devices are CD writers, DVD devices, and zip drives.

Figure 5-5: Looking at the IDE/ATA connectors on the motherboard.



Ultra DMA

Ultra DMA (Direct Memory Access) drives have two major benefits over ATA drives:

- ◆ **Speed:** Originally, Ultra DMA devices functioned at twice the speed of regular IDE devices. IDE devices execute commands at 16.6 MBps, whereas Ultra DMA devices execute commands at 33.3 MBps, or 66.6 MBps with Ultra DMA/66. Note that there have been improvements to Ultra DMA to include speeds of 133 MBps.
- ◆ **Reliability:** Ultra DMA devices implement error correction, which provides for increased data reliability compared with IDE, which does not implement error correction.

To take advantage of Ultra DMA technology, you need an Ultra DMA drive and an Ultra DMA-compatible BIOS. In addition, you need an Ultra DMA-compatible driver loaded in the operating system that uses the device.

It is important to note that Ultra DMA technology is backward-compatible with IDE and EIDE. For example, if you have a motherboard with Ultra DMA support, you can still plug an IDE or EIDE device into the controllers on the motherboard. You can also install an Ultra DMA drive on an EIDE board.

Ultra DMA/66 requires the use of an 80-wire ribbon cable, which contains the normal 40 wires of an IDE cable plus an additional 40 ground wires to reduce noise on the cable and thus increase performance. The 80-wire cable still uses a 40-pin connector on the end to keep compatibility with IDE.

Drive performance

The *performance* of the drive is determined by the *Programmed I/O* (PIO) mode or DMA mode it supports. The PIO mode or DMA mode is a protocol that determines the transfer rate of the drive. A number of different PIO modes and DMA modes are supported by devices:

- ◆ **PIO Mode 0** runs at 3.3 MBps and is used by ATA-1.
- ◆ **PIO Mode 1** runs at 5.2 MBps and is used by ATA-1.
- ◆ **PIO Mode 2** runs at 8.3 MBps and is used by ATA-1.
- ◆ **PIO Mode 3** runs at 11.1 MBps and is used by ATA-1 and ATA-2.
- ◆ **PIO Mode 4** runs at 16.6 MBps and is used by ATA-1 and ATA-2.
- ◆ **DMA 33** runs at 33.3 MBps and is used by Ultra DMA drives.
- ◆ **DMA 66** runs at 66 MBps and is used by Ultra DMA drives.
- ◆ **DMA 100** runs at 100 MBps and is used by Ultra DMA drives.
- ◆ **DMA 133** runs at 133 MBps and is used by Ultra DMA drives.



For the exam, be sure to know what a PIO mode is and also the transfer rates of Ultra DMA drives.

Installing IDE devices

In this section, I show you how to install an IDE device. Because all hard drives are different, I present only the most popular solutions for installing hard drives.

First, you want to open up the casing of the computer and find an empty bay to mount the new hard disk into. If you are removing the old hard drive, you will be able to use the same bay. However, if you are adding a hard drive, you have to find an empty bay and mount the drive in place. Second, you slide the drive in the bay so that the power connector and ribbon cable connector are facing the motherboard. Figure 5-6 shows an IDE drive being placed into the mounting bay.

IDE cabling

After you have mounted the drive in place, as a third step, connect the IDE connector on the motherboard to the IDE connector on the hard disk with a 40-wire IDE ribbon cable and then connect a power connector to the drive. Figure 5-7 shows an IDE ribbon cable.



Figure 5-6:
Mounting
the hard
disk into the
drive bay.




Figure 5-7:
Looking at
an IDE
ribbon
cable.



For the exam, remember that an IDE hard drive uses a 40-wire ribbon cable, while a floppy drive uses a 34-wire ribbon cable. Also note that Ultra DMA uses an 80-wire ribbon cable that contains 40 wires for data and 40 additional grounding wires.

One of the wires (known as Wire 1) on the IDE ribbon cable is a different color than the others; usually, it's red, but it may be blue. Wire 1 must be placed over Pin 1 when connecting the ribbon cable to the hard drive and motherboard — a procedure known as the *Pin 1 rule*.

The big question is: How do you know which pin is Pin 1? Hopefully, the manufacturer has indicated Pin 1 by placing a small “1” near it. If you look at the connector on the hard drive and on the motherboard, you may see a small “1” on one end of the connector. That’s where you need to place Wire 1 when connecting the ribbon cable. Sometimes the manufacturer will place a “40” by Pin 40 instead of displaying where Pin 1 is, so watch for that as well.



Sometimes the manufacturer may not have enough space to indicate Pin 1 and so does the opposite, which is to indicate Pin 40. This method gives you the same information, though: If you know what side Pin 40 is on, you know that Pin 1 is on the opposite side.

After you have connected the IDE ribbon cable, you want to give the hard disk power from the power supply so that you can run the motor in the drive. (People often forget this step and then wonder why the drive doesn’t work.) Figure 5-8 shows how to connect the ribbon cable and the power supply cable to the hard disk.




Figure 5-8: Connecting an IDE drive.

Master/slave configuration

If you are installing multiple IDE devices, you are creating an IDE chain. The chain will be made up of one 40-wire IDE ribbon cable with two drives connected to it. Most IDE ribbon cables today have three IDE connectors on them — one that connects to the motherboard and one for each of the two drives that can be connected in a single chain.

After you have the two drives connected together, you need to configure the drives into a master/slave configuration. Why this type of configuration? Because each drive has a built-in controller that makes it act like its own boss. Have you ever tried to work in an environment with two bosses? The purpose of designating a master is to specify who the boss of that chain is — the controller that will be responsible for communicating with the processor.

When setting up a multi-drive system, you have two drives, each with a controller that can potentially send and receive signals to and from the processor. To save confusion, one of the drives is designated as the *master*. The master drive receives all signals from the processor and sends back any data on behalf of both drives. The other drive is designated as the *slave*. The slave drive passes any information it wants to send to the processor up to the master, which then forwards that information to the processor. Figure 5-9 illustrates a multi-drive system.




Figure 5-9: One drive is configured as the master and the other is configured as the slave.

To configure the master/slave relationship, you have to configure the jumper settings on the drives. The idea of the procedure is similar for all drives, but the actual jumper setting may be different for each type of drive — the best thing to do is review the documentation for the drive.

Most hard drives today have the jumpers labeled as “Master” and “Slave.” Configuring these drives is usually pretty straightforward. You place the jumper on the master setting for the drive you would like to be the master of the chain — meaning this drive will communicate with the processor directly — and you place the jumper to the slave setting for the drive you would like to configure as the slave drive. The slave drive receives instructions from the master. Figure 5-10 displays master and slave jumper settings.




Figure 5-10: Adjust the jumpers to set one drive as a master and the other as a slave.

Many times, configuring your drives in a master/slave setup may not be as easy as it looks. What do you do when you cannot see a specific master or slave jumper setting? Many manufacturers place documentation on the back of the drive telling you how to configure the drive for master and slave setups. Unfortunately, this was not the case on a drive I had to configure in the office one day. So I went to the manufacturer's Web site and found out that this drive had a jumper labeled J20; the Web site documentation stated that if the drive was the only drive or was the master in a dual-drive configuration, the jumper (J20) should be set. If the drive was the slave in the dual-drive configuration, the jumper (J20) should be removed. Figure 5-11 shows the configuration for this example.

The moral of the story: If you can't find the jumper settings on the drives themselves, check out the drives' documentation or go to the manufacturers' Web sites.

Today, you are probably working with EIDE drives which, as mentioned earlier, support up to four devices instead of two. These systems have two IDE connectors on the motherboard: the primary IDE connector and the secondary IDE connector. Each connector can have two devices configured in a master/slave relationship, for a total of four devices. Typically, the devices can be hard drives, CD-ROMs, or DVD devices. The steps are the same for configuring master and slave relationships for a hard disk, CD-ROM, and DVD because they are all considered IDE devices. Normally you configure the CD-ROM or DVD as the slave, but you could also make it the master device on the second IDE connector.

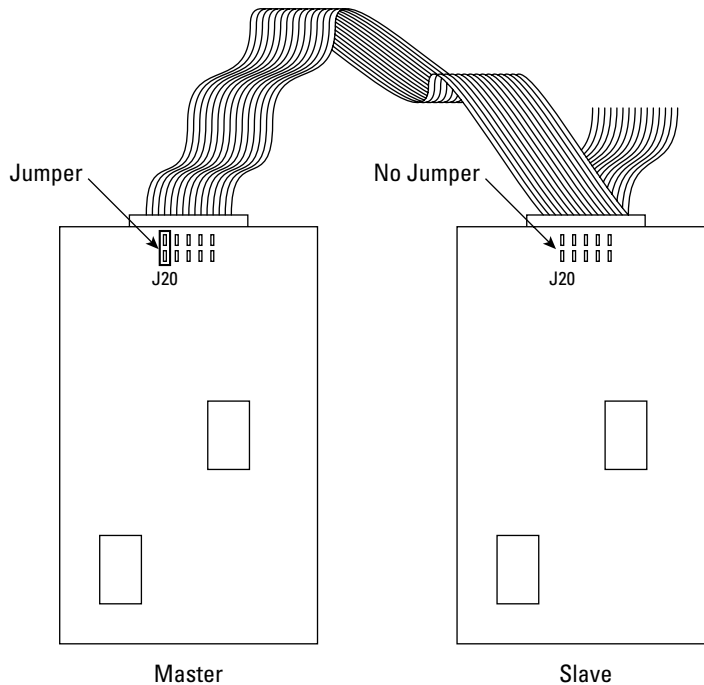


Figure 5-11: Looking at an additional jumper configuration for a dual-drive setup.

After configuring the jumpers on the drive, you power on the system. The computer should recognize that there is an additional drive and take you into the CMOS setup program. When you review the CMOS settings, you should notice the newly added drive. If you don't see the newly added drive, verify the master and slave settings and maybe even switch the roles of the drives around and try again.



Lab 5-1 demonstrates how to master and slave the drives. Lab 5-1 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Learning How SCSI Works

In the following sections, I discuss a type of device that is more popular than IDE drives in high-end machines, such as powerful workstations or servers — SCSI devices.

SCSI overview

SCSI is an acronym for *Small Computer System Interface*. The important part of this term is *small computer*, meaning that SCSI has its own brain, known as the *SCSI adapter*, that handles the SCSI environment. This *SCSI adapter* (also

known as a *SCSI card* or *SCSI controller*) is responsible for managing all SCSI devices and controlling the conversation on the SCSI chain.

SCSI technology has many advantages over IDE technology, such as

- ◆ **The types of devices supported:** SCSI supports a multitude of devices, including hard drives, CD-ROMs, scanners, printers, and tape drives, to name just a few. This is a huge benefit because originally IDE typically only supported hard drives and CD-ROMs.
- ◆ **The number of devices supported in a single SCSI chain (also known as a SCSI bus):** Original versions of SCSI supported up to eight devices in the chain, but one of those devices is the SCSI card that's added to the computer to give you the capability to use SCSI.

Remember that IDE only allows two devices in the chain, and EIDE supports four devices, so with SCSI you are not only allowed to have more types of devices, but you are also allowed to have more of those devices!

Original SCSI supports up to eight devices in the chain, but if the exam asks how many devices can be *attached* to a SCSI adapter, the answer is seven. If the test asks how many devices can *exist* in the SCSI chain, the answer is eight — the card counts as one of those devices. Remember to watch the wording of the questions closely.

- ◆ **The performance of SCSI over IDE devices:** Original SCSI devices don't compare in the performance category with EIDE devices, but some of the later SCSI technologies, such as SCSI-2 and SCSI-3, can outperform IDE and EIDE.

I give details about transfer rates of the different types of SCSI devices in the “Types of SCSI” section, later in this chapter.

For the remainder of this chapter, assume that you are dealing with original SCSI, also known as SCSI-1. In the “Types of SCSI” section, I discuss the newer types of SCSI.



Host adapter

When you install SCSI devices, you first need to install the SCSI host adapter. The *SCSI host adapter* is an expansion card that you add to the computer so you can chain SCSI devices off the adapter. In essence, the SCSI host adapter is the brain of the SCSI bus; it acts as the controller for the SCSI bus. Figure 5-12 shows the SCSI adapter being inserted into an expansion slot in the system.

The SCSI controller (adapter) is responsible for sending and receiving all information to and from the SCSI bus, just like the IDE controller. When the system has information for one of the devices in the SCSI bus, the system hands the information over to the SCSI controller, which then passes the

information to the appropriate device in the chain. Figure 5-13 shows a SCSI bus — made up of the SCSI host adapter (the SCSI controller) — along with two internal SCSI devices and two external SCSI devices.




Figure 5-12:
The SCSI
host
adapter.

The beauty of the entire setup is that the SCSI adapter in the computer is assigned resources, such as an IRQ and an I/O address. Each device in the SCSI chain is not assigned these resources because all processor information passes to the SCSI controller, and the controller passes the information to the devices. This means that the system never talks to the devices directly, so each device does not require an IRQ and an I/O address.

When you go out to purchase a SCSI adapter, you first have to look inside your system to figure out what type of expansion slots are free. Today, you will typically have some PCI slots, but you may have an ISA or an EISA slot, as well. The difference between these expansion slots is performance — PCI runs at 33 MHz, while ISA and EISA run at only 8 MHz. Also, PCI and EISA are 32-bit technologies, while ISA is only 16-bit. The bottom line is that if you have some PCI slots free, you will probably end up purchasing a PCI SCSI adapter.

So the next big question is when the SCSI controller receives information for a particular device in the chain, how does it send the information to that device?




Figure 5-13:
Looking at a
SCSI daisy
chain.

Addressing

Each device is assigned an internal address, a *SCSI ID*, in the SCSI bus. The SCSI controller knows the address of each device. When the SCSI controller receives information for a particular device, the controller references that device by its ID in the SCSI bus. This way, there's no confusion as to whom the data is destined for.

You are responsible for assigning the SCSI IDs when you connect each device to the SCSI chain. You assign an ID either by jumpers or DIP (dual inline package) switches if the device is an internal device, or by a spinner if the device is external. A spinner is an indicator on the back of the external SCSI device whose value you can change by pressing the button to increase or decrease the SCSI ID. Figure 5-14 shows the back of an external SCSI tape drive and how to change the SCSI ID using the spinner.




Figure 5-14:
Assigning a
SCSI ID to
an external
device.



In Figure 5-14, you can also see the type of connector used for external SCSI devices. This is a Centronics 50-pin connector.

If you are installing an internal SCSI device, you will most likely need to assign the SCSI ID by using jumpers. Internal SCSI devices have a jumper set with three pairs of jumper pins. The decimal values of these jumpers, although probably not shown on the drive, are 4, 2, and 1 (from left to right). Table 5-4 shows this jumper setup.

Table 5-4 Jumper Block Settings

	<i>J2</i>	<i>J1</i>	<i>J0</i>
Decimal Value	4	2	1
SCSI ID 0	0	0	0
SCSI ID 3	0	1	1
SCSI ID 6	1	1	0

If you would like to assign the device a SCSI ID of 0, then you would not jumper any of the jumper pins. This is shown in Table 5-4 by having the off state (represented by 0) at each jumper location. Another example of setting a jumper ID would be if you wanted to assign the device a SCSI ID of 3, you would jumper the *J0* pins and the *J1* pins, but not the *J2* pins. This would enable decimal values of 1 + 2. Table 5-4 also shows what would happen if you wanted to set the SCSI ID to 6. You would enable jumpers *J1* and *J2*, which enable the decimal values of 2 + 4, while the jumper *J0* has an off state.

The IDs you assign to each device are completely up to you, but note that the higher the number, the more important the device is to the SCSI bus. For example, if two devices need to send information through the bus at the same time, the device with the higher SCSI ID is always given priority. This is why the SCSI host adapter is usually assigned the highest number in the chain (usually 7 if the bus supports eight devices — the ID numbers start with 0).


So the host adapter is assigned an ID of 7, and a SCSI bootable hard drive is assigned a SCSI ID of 0, which is the SCSI ID that the controller automatically looks to in order to boot off a SCSI hard drive. Outside of that, you decide what the ID numbers are for each device. As a rule of thumb, give slower devices higher ID numbers so that they receive a higher priority in the SCSI bus. Devices with higher priority get access to the bus before lower priority devices.

SCSI cabling


Different types of cabling are used to chain SCSI devices to the SCSI adapter. Internal devices use a 50-wire ribbon cable; external devices require a thick Centronics cable to connect to the Centronics 50 (typically used by SCSI-1) or the Centronics 68 (typically used by wide SCSI-2 technologies or Ultra SCSI-3) connector on the back of the device.




The different versions of SCSI use a large number of different cable types. Figure 5-15 shows some internal SCSI connectors, and Figure 5-16 displays a handful of external SCSI connectors. Be sure to be familiar with these cable types for the exam.




50-pin IDC female



50-pin IDC male



68-pin high density - male



80-pin SCSI SCA connector

Figure 5-15:
Looking at
internal
SCSI
connectors.





Figure 5-16:
Looking at
external
SCSI
connectors.

Termination

Both ends of the SCSI bus must be terminated so that when a signal is sent down the SCSI bus, it is absorbed at the end of the bus by the terminator. If the signal was not absorbed, or removed from the bus with a terminator, the signal would bounce back and collide with other data on the bus. A collision would destroy the signal. The first device in the chain must be terminated along with the last device in the chain, the first device usually being the host adapter. If the device is an internal device, terminating may involve modifying jumper settings. If the device is external, a terminator will be added to the back of the device. Figure 5-17 shows a terminator for external devices.

Figure 5-17:
Looking at a
SCSI
terminator.



If your SCSI chain has a combination of internal and external SCSI devices, then the card shouldn't be terminated because it is no longer the end of the SCSI chain. Instead, you should terminate the devices at either end of the SCSI chain.



Most SCSI cards today are self-terminating so you typically will not need to terminate them manually yourself.

Types of SCSI

Over the last twelve years or so, SCSI technology has increased in performance to stay competitive with advances in IDE and EIDE. Newer versions of SCSI have amazing transfer rates, which is one of the reasons why you find network servers using SCSI hard drives instead of EIDE devices. The following list outlines the key points about the different versions of SCSI:


- ◆ **SCSI-1:** The original version of SCSI, SCSI-1, was an 8-bit technology with a transfer rate of 5 MBps. One of the major benefits of SCSI was that you weren't limited to two devices in a chain like you are with IDE. SCSI-1 allowed you to have eight devices in the chain, with the controller counting as one.
- ◆ **Fast SCSI-2:** Fast SCSI-2 increased the performance of SCSI by doubling the transfer rate. Fast SCSI-2 devices transfer information at 10 MBps. Fast SCSI-2 is still an 8-bit technology and supports eight devices in the chain.
- ◆ **Wide SCSI-2:** Wide SCSI-2 doubled the 8-bit data path of SCSI to 16 bits. Doubling the width of the data path raised the transfer rate to 10 MBps, like Fast SCSI-2, but Wide SCSI-2 can support 16 devices in a chain.

When trying to remember the difference between SCSI-1, Fast SCSI-2, and Wide SCSI-2, think of it this way: *fast* implies speed, so the transfer rate is increased. *Wide* implies "wider" or bigger, which is the data path that has been increased; as a result, you also get a higher transfer rate.

- ◆ **Fast Wide SCSI-2:** Fast Wide SCSI-2 is the combination of Fast SCSI-2 and Wide SCSI-2. The data path of Fast Wide SCSI-2 is 16 bits, the transfer rate is 20 MBps, and it supports 16 devices in a chain.



- ◆ **Ultra SCSI:** Ultra SCSI takes the transfer rate of 10 MBps and doubles it again to 20 MBps! With Ultra SCSI, the bus width is only eight bits, and the number of devices that exist in the chain is eight.
- ◆ **Ultra Wide SCSI:** Ultra Wide SCSI is Ultra SCSI with the bus width increased to 16 bits, and the number of devices in the chain is increased to 16. The transfer rate of Ultra Wide SCSI increased to 40 MBps.
- ◆ **LVD (Ultra2 SCSI):** *Low Voltage Differential (LVD)*, also known as *Ultra2 SCSI*, has a bus width of 16 bits and supports up to 16 devices. LVD is a popular SCSI version due to having a high transfer rate of 80 MBps.



Be prepared to answer questions on the A+ exam about the different types of SCSI. I suggest memorizing the transfer rates of each type.

Installing SCSI devices

If you understand the issues with SCSI, installing a SCSI bus is fairly simple. First, you want to assign a unique ID number to each device. I usually perform this step at the beginning so that when everything is connected, I won't have to play around figuring out how to change the ID of the devices. For more on assigning the IDs, see the section, "Addressing," earlier in this chapter.

When assigning the ID numbers, remember that you want to assign the bootable drive the ID of 0 because the SCSI host adapter automatically looks to SCSI ID 0 for a bootable device. You also want to enable the SCSI BIOS on the SCSI controller if you are booting off a SCSI hard disk. When you enable the SCSI BIOS, you won't need to install a driver for the card because the PC will recognize the device on startup. If you are booting off an IDE drive and using the SCSI disk as an additional drive, you should disable the SCSI BIOS and install a driver in the operating system.

After you assign the ID numbers to each device, insert the SCSI host adapter into the expansion slot of the PC. After inserting the SCSI card, chain all of the devices together.



When preparing for the A+ exam, it is important to know the different terms used for a particular technology. Another term for a *SCSI chain* is a *daisy chain*. When you chain the devices together, you are creating a daisy chain.

When you have the devices chained together, make sure that each end on the SCSI bus is terminated. If the last device is an external device, you need to put the terminator on the end of the device. If you are installing internal devices, you need to check the documentation on the internal devices to find out what jumpers to set.

At this point, the IDs are configured for each device, the SCSI card is inserted into the PC, and the devices are connected to the card to create a SCSI bus.

We also terminated the SCSI bus at either end. Before installing the driver for the SCSI card in the operating system, I'll review the steps to install a SCSI device one last time. To install a SCSI bus:

- 1. Assign unique IDs to each device.**
- 2. Install the SCSI host adapter into the expansion slots.**
- 3. Chain devices to the SCSI host adapter.**
- 4. Terminate the SCSI bus at both ends of the chain.**
- 5. Install the driver for the SCSI card if you are not booting off the first hard disk.**

After you have connected the hardware for the SCSI chain, you need to load a driver in the operating system for the SCSI host adapter. After inserting the SCSI card and powering on the Windows operating system, Plug and Play should kick in, asking for the driver of the new hardware. If Plug and Play does not kick in, you can run the Add Hardware Wizard found in the Control Panel or My Computer properties of the system. To install a SCSI adapter in Windows 2000/XP, and Windows Server 2003, follow these steps:

- 1. Choose Start⇒Settings⇒Control Panel in Windows 2000 or right-click My Computer and choose Properties in Windows XP and Windows Server 2003.**
- 2. Double-click the Add Hardware icon in Windows 2000. If you are using Windows XP, go to the Hardware tab and click the Add Hardware Wizard button.**

For both methods, the Add Hardware Wizard now displays a welcome screen.

- 3. Click Next.**
- 4. Select Add or Troubleshoot a Device and then click Next.**
- 5. In the Choose Hardware Device dialog box, select Add a New Device and then click Next.**
- 6. For these steps, you will choose the adapter to install, so select No, I Want to Select the Hardware from a List. Click Next.**

If you know the host adapter you are installing, select No, I Want to Select the Hardware from a List. If you are unsure of the host adapter you are installing, then choose Yes, Search for New Hardware.

The next screen asks you what type of device you are installing.

- 7. Choose SCSI and RAID controllers and then click Next.**
- 8. In the list of devices, select the manufacturer on the left side and the adapter to install on the right side and click Next.**

9. Click Finish.

10. Restart the computer.



To practice installing SCSI devices in Windows perform Lab 5-2. Lab 5-2 can be found the Labs .pdf file in the Author directory of the CD-ROM.

The Serial ATA Hype!

IDE technology has been around for many, many years, and there has been a big need for a change in hard drive technology — that change came as a new hard drive interface called *Serial ATA (SATA)*. IDE is a parallel technology, and though SATA is a serial technology, it offers great speed and other benefits. SATA is also a lot faster than IDE — approximately 30 times faster, with current speeds of 150 Mbps and future speeds of 600 Mbps.

One of the first benefits of SATA is that it is a *hot-swappable* technology, meaning that you can add or remove drives from the system without shutting the system down. This is a huge benefit when you look for RAID solutions for servers that you don't want to spend a lot of money on — like a server for a small company. (For more on RAID, see the section “Securing Data with RAID,” later in this chapter.)

Another benefit of SATA is in the cabling. Because SATA is a serial technology, the cables can be longer than your typical IDE ribbon cables. I don't know how many times with IDE I had to switch the CD-ROM and my second hard drive around just so the ribbon cable could reach. SATA cables can be 39 inches long, while the maximum distance for IDE is 18 inches.


The other benefit of the cabling with SATA is that it uses only 7 wires, as opposed to the 80 wires used in newer IDE drives. The benefit here is that it allows for better airflow in the system, which results in a cooler system. Figure 5-18 shows a SATA drive and cable.



If you don't have a motherboard that has SATA connectors on it, then you can get a PCI card that does have the connectors. Also, for backward compatibility, you can get a SATA bridge that allows IDE drives to be connected to a SATA system.

With IDE, you have to master and slave the drives when you want multiple IDE devices. Yet another benefit to SATA is that there is no mastering and slaving of devices because you can only connect one SATA device to a single connector, or channel. The reason this is a benefit is because when installing a drive you don't need to have knowledge of mastering and slaving devices — you simply connect the SATA device to the system!

Figure 5-18: A SATA cable and power connector connecting to a SATA drive.



SATA power connector

SATA data connector

Book II
Chapter 5

Working with
Storage



If you want two SATA drives, then you need two SATA connectors — and don't forget that you can get cards to give you more connectors allowing you to install additional SATA devices.

Using Removable Storage

Not all data storage is hard-wired to the system. *Removable storage* is the term used for storage media that you take away from the system with you.

Floppy disks

The floppy disk has been very popular in the past for storing data and being able to carry the data with you. Floppy disks are not as popular today because they typically can store only up to 1.44MB of data. Today, one picture from a digital camera will use that space up.

A floppy disk has a hard plastic casing with a sliding metal shutter that allows the drive to access the silicon disk inside. The floppy drive comes in different sizes, or formats:

- ◆ **5¼-inch:** The 5¼-inch floppy disk came in two formats: 360KB and 1.2MB. The 5¼-inch floppy disk had a soft, flexible shell that did not protect the disk very much. You may not encounter a 5¼-inch floppy because they have been obsolete for many years.
- ◆ **3½-inch:** A 3½-inch floppy disk is the floppy disk that you will most likely encounter, if you encounter one at all. As with the 5¼-inch floppy disk, there is a protective shell but the shell of a 3½-inch floppy is a hard plastic shell. The 3½-inch floppy comes in two formats: 720KB and 1.44MB. The 1.44MB format is more common.

CD-ROM/CDRW

A *Compact Disc–Read Only Memory (CD-ROM)* is an optical storage technology that uses a laser to read and write data. Originally, as the name implies, one could only read from CD-ROMs. CD-ROMs are the preferred media to distribute software and fairly large amounts of data.

Originally, CD-ROMs stored 650MB of data and could store 74 minutes of music, but today's CD-ROMs store 700MB of data or 80 minutes of music. CD-ROMs are written to from the inside out, and if you look at the bottom of the CD-ROM, you will see the lines indicating where data was written — very useful information if you ever pick up a CD and wonder if it was written to.

The speed (transfer rate) of the CD-ROM is measured in multiples of 150 KBps and is indicated with an *X*. For example, an old 1X (pronounced “one times”) CD-ROM had a transfer rate of 150 KBps, while an 8X has a transfer rate of 1200 KBps ($150 \times 8 = 1200$), and a much faster 48X has a transfer rate of 7200 KBps.

Today, CD drives are writeable: if you want to *burn* your own CD you can. (Writing to a CD is often called *burning a CD*.) This makes CD drives much more popular than the older floppy drives due to the amount of information you can store on the CD — a great solution for backing up your data!

There are two types of writeable CDs, a *CD Recordable (CDR)* and a *CD ReWritable (CDRW)*. The difference between a CDR and CDRW is that you can write to the CDR only once, while you can reuse the CDRWs many times by reformatting the CD and starting again.

A CDR is an example of a *Write Once Read Many (WORM)* disc.

You can write to a CDR multiple times, but you cannot overwrite areas of the disc that have already been written to. This means that with a CDR if you write to it many times the additional write operations are appending the information to the end of the CDR. Each burn operation that you perform is called a *session*, and most writeable CD drives today support multiple sessions. For example, say you back up your pictures to CD and use only 350MB



of space on the CDR. You can write more data to the CD with the remaining space at a later time. With the cost of CDR being so low nowadays, I typically don't bother. I burn a CD, label it, and then file it away.

When writing to a CD, the process is not done magnetically like it is with hard drives — as mentioned earlier, the write operation is performed with a laser. The CD has a chemical-dye layer mixed with a thin reflective layer. When you write to the CD, the chemical layer is heated with the laser to create an solid state at that location. These locations reflect less light, and the different patterns of reflection create the data on the disk.

DVD/DVDRW

Digital Versatile Disks (DVDs) are similar to CDs in the sense that they are another type of optical storage — but they store a lot more data. The typical DVD stores 4.7GB of data. Some DVDs (unlike CDs) can store data on both sides of the disk, and newer DVDs even store data on different layers on the DVD. This allows the DVD to store more than the 4.7GB, depending on the DVD standard. Table 5-5 describes the different DVD standards.

<i>Standard</i>	<i>Specifications</i>	<i>Total Storage Space</i>
DVD-5	Single-sided, single layer	4.7GB
DVD-9	Single-sided, double layer	8.5GB
DVD-10	Double-sided, single layer	9.4GB
DVD-18	Double-sided, double layer	17GB

Flash drives

Flash drives, also known as *thumb drives* or *memory sticks*, are the popular method for carrying data from computer to computer. Flash drives can store anywhere from 64MB to multiple gigabytes of information and are relatively cheap.

Flash drives are USB devices that you simply plug into the USB port on the computer. Plug and Play kicks in, detects the device, and assigns the drive a letter in the My Computer icon. To access the flash drive, you simply double-click the drive in My Computer and open, copy, and move files as you wish. Figure 5-19 shows a flash drive.

USB external drives

External drives are just as popular as flash drives today. Like flash drives, an external drive uses a USB connection. Unlike flash drives, external drives allow you to store hundreds of gigabytes, even a terabyte, of data. These are great solutions to add more space to a laptop computer.

Figure 5-19:
A flash drive uses a USB connection to interface with the computer.



Zip drives

A zip drive is similar to a floppy drive, but it is a little bit larger physically and stores more information. A zip drive can store 100MB or 250MB of data, depending on what type of drive it is. These drives were popular solutions for backing up data before CD-ROMs and external drives became popular.

Flash cards

Multimedia devices, such as MP3 players and digital cameras, now support adding memory to the device by using flash cards (shown in Figure 5-20). The benefit of flash cards is that most computers and laptops have ports on them that you can insert the memory card into. This allows you to upload pictures or music to the memory without needing to connect the camera, which requires connecting a USB cable to the computer.

Figure 5-20:
A Compact-Flash card used to store data on a digital camera.



Tape

A popular type of media for data is *tape*, which is typically used to store backup copies of the data. This copy of the data stored on tape is used to bring the data back if the hard drive fails. Different types of tape drives are popular today:

- ◆ **Quarter-Inch Cartridge (QIC):** QIC is one of the oldest and most popular tape standards that supports many different tape sizes. The QIC size depends on the standard, and each standard is labeled similar to QIC-80 — the QIC tape that supports 80MB of data. Other examples are QIC-40 (40MB tapes) and QIC-5210 (25GB tapes).
- ◆ **Travan:** Travan is a tape drive standard that is based on the QIC standard but supports compression. Travan is typically more expensive and can store from 400MB to 8GB of data or more. Examples of tape sizes in this standard are labeled as TR-1, which stores 400MB of data, and TR-4, which stores 8GB of data.
- ◆ **Digital Audio Tape (DAT):** DAT drives store data in a digital format and use two heads, one for reading data and the other for writing data. DAT drives use a standard called *Digital Data Storage (DDS)* to store data. Examples of DDS tapes are DDS-1, which can store 2GB of data, DDS-2, which can store 8GB of data, and DDS-4, which can store 40GB of data.

Understanding File Systems

The file system dictates how information is organized on the disk. For example, the file system determines how large the *allocation unit*, or *storage unit*, of a file is. If you create a 12K file, how much space is that file really using — 12K, 16K, or 32K? Such organizational issues are what the file system deals with.

The following sections introduce the different file systems available and the operating systems that support them. You also find out what features the file systems do and do not support.

The FAT file system

The *File Allocation Table (FAT)* file system has been the most popular file system up until the last few years. Although the FAT file system is the most common (it can be used by all operating systems), it is losing the popularity contest to its successor — FAT32 — due to its age and limitations.

The FAT file system was the file system used by DOS, Windows 3.1, and Windows 95, and is supported in Windows NT, Windows 98/Me, and Windows 2000/XP/2003. FAT's biggest strength is that it's the file system most widely understood by different operating systems — but it has many shortcomings. One of the major shortcomings is that it cannot create a partition larger than 2GB. (A discussion of partitions is coming up in the “Managing Partitions and Volumes” section; for now, consider a *partition* simply as a discrete portion of space on the disk.) The 2GB size limit was not a major limitation until hard drive sizes exceeded a few gigabytes. For example, a problem with the FAT file system is that a 20GB drive would need to be divided into 10 partitions to use all the space — an impractical and inappropriate use of space. Can you imagine being required to divide your home up into ten different

rooms whether you wanted to or not? Instead of five spacious rooms, you'd get ten cramped rooms! Not practical!

Earlier in this chapter, in the section “Disk geometry,” you discover the characteristics of a disk, including *clusters*, which are groups of sectors (each sector taking up 512 bytes on the disk). To refresh your memory, a cluster is what a file is written to, and only one file can occupy a cluster at a time. The cluster size is determined by the partition size and the file system being used. For example, you may have a 2GB FAT partition with a 32K cluster size. The issue with clusters is that if you have a 32K cluster size and you save a 12K file, then you waste 20K of hard disk space because only one file can occupy a cluster. Over time, as more files are saved, this could add up to a lot of wasted space! The solution is to use smaller partition sizes, which create smaller cluster sizes, or to use a different file system that uses smaller cluster sizes.

Bottom line, the FAT file system uses clusters inefficiently. Table 5-6 lists the cluster sizes used with different partition sizes on FAT file systems.

<i>Partition Size</i>	<i>Cluster Size</i>
0–127MB	2K
128–255MB	4K
256–511MB	8K
512–1023MB	16K
1024–2048MB	32K

One other limitation of the FAT file system is its lack of built-in security. For example, you may want to set permissions on a folder so that only Bob can create files in that folder. Unfortunately, you can't set permissions on a FAT file system. In fact, if you are using a FAT file system with Windows 2000 or Windows XP (which are well respected for their security), you will lose the capability of securing the file system even though you are told you are running a much more secure system than a Windows 9x system. Bottom line, the file system dictates the feature set you get when dealing with files and folders, so no matter what OS you are running, if you are using the FAT file system, you can't secure the files.

So, besides being a file system that all operating systems can run on, what is the FAT file system good for? Well, the versatility of FAT may be its only major benefit — most operating systems can run on it, which makes it a great file system for configuring a dual-boot system. A *dual-boot system* has multiple operating systems installed and has the capability to boot to either operating system at any point in time. For example, if you were supporting an organization that was dual-booting its systems with Windows NT and Windows 98, you would need to use a file system that was common to both operating systems — and the FAT file system is common to all operating systems.

Reasons to dual boot

There are a number of reasons you may want to have a dual-boot system:

- ✓ Software doesn't run smoothly on one operating system, so you install an additional operating system and boot to that OS to run your software.
- ✓ You are part of the help desk team in an organization that runs two different operating systems, and you will need to boot to the appropriate OS to find solutions to problems.
- ✓ You are studying for the A+ exam and need to run Windows 2000 and Windows XP but don't have two different computers.

The FAT32 file system

After the retail release of Windows 95, an update to the operating system was created, known as Windows 95 OSR2 (OEM Service Release 2). Windows 95 OSR2 introduced an updated FAT file system called *FAT32*. One of the apparent benefits of FAT32 was that the maximum partition size was increased from 2GB to 2 *terabytes* (TB). Now, when you go out and buy that 20GB drive, you don't have to divide it into 10 partitions; you can keep one 20GB partition.



Although FAT32 has the capability to have partitions of 2 TB in size, Microsoft has limited the size of FAT32 partitions in Windows 2000/XP operating systems to 32GB — their reasoning for this is that you should be using NTFS as a file system.

The other major benefit to FAT32 is that it dramatically decreases the cluster size to make better use of disk space. Table 5-7 compares the default cluster sizes of FAT partitions with the default cluster sizes of FAT32 partitions. Note that FAT32 doesn't support partitions smaller than 512MB.

Table 5-7 Comparing Cluster Sizes

<i>Partition Size</i>	<i>FAT Cluster Size</i>	<i>FAT32 Cluster Size</i>
0MB–127MB	2K	Not Supported
128MB–255MB	4K	Not Supported
256MB–511MB	8K	Not Supported
512MB–1023MB	16K	4K
1GB–2GB	32K	4K
2GB–8GB	Not Supported	4K
8GB–16GB	Not Supported	8K
16GB–32GB	Not Supported	16K
Over 32GB	Not Supported	32K

Looking at Table 5-7, if you have a 2GB partition on FAT32, you have a cluster size of 4K, but if you had created the partition and used the FAT file system, the cluster size would be 32K. If you save a 12K file on the FAT partition, you lose 20K of disk space, while the same 12K file would waste no space on a FAT32 partition because it will span three clusters of 4K each!

One of the shortfalls of FAT32 is that MS-DOS, Windows 3.1, the original release of Windows 95, and Windows NT 4.0 do not support it. This means that only Windows 95 OSR2, Windows 98, Windows ME, and Windows 2000/XP/2003 can access data on a FAT32 partition, which is great considering these are the popular operating systems today.

NTFS

Starting with Windows NT, Microsoft implemented a new file system called *New Technology File System (NTFS)*. NTFS makes better use of the space available on a particular disk by using 512 bytes as the cluster size, which is the same size as a sector! This means that you are wasting even less space on an NTFS file system than on a FAT32 file system.

The original version of NTFS supported a number of features that made it more attractive than the FAT versions of the file systems. With NTFS, you could configure permissions that controlled who could access what files. You could also take advantage of features such as compression and auditing.

One of the biggest complaints with the original version of NTFS is that it had no way to limit how much disk space a user could use. As a result, users could waste gigabytes of hard disk space on the server, and the administrator could not stop the user unless a third-party program was purchased. Limiting disk space usage is one of the improvements that Microsoft made on the next version of NTFS, known as *NTFS version 5.0*, which was implemented with Windows 2000 and every Windows OS after that.

NTFS 5.0

This newer version of NTFS has a few extra features over original implementations of NTFS, one of which has been long overdue — *disk quotas*. Disk quotas allow the system's administrator to choose the amount of disk space that each user is allowed to use by placing a limit on the disk. For example, when managing the home directories, you can ensure that Bob is not allowed to use more than 500MB of disk space.

Another feature of NTFS 5.0 is the *Encrypting File System (EFS)*. EFS uses public key/private key technology to encrypt a file stored on the hard drive. When a file or folder is encrypted with EFS, only the person who created the file or the *recovery agent* (by default, the administrator is the recovery agent) can open the file. When using EFS, even if another user has permission to view the file, he or she will be unable to do so because the file is

encrypted. The encrypting file system is a big selling point for organizations with mobile users who need to protect the privacy of the data that sits on their laptops.

To summarize, the NTFS file system offers the following features over FAT and FAT32:

- ◆ **The capability to secure the resource through permissions**
- ◆ **The capability to secure files through encryption**
- ◆ **The capability to enable auditing to monitor who accesses the files and folders**
- ◆ **The ability to compress the contents of a file or folder**

HPFS

Years ago, the *High Performance File System (HPFS)*, which gained its popularity with the OS/2 operating system, was a major improvement over the FAT file system. Some of the benefits of OS/2 are that it supports long file names, up to 254 characters (including the path). HPFS also supports partition sizes up to 2000GB and uses a cluster size of 512 bytes! When looking at the benefits of HPFS, you may be thinking, “What’s the big deal, I get that with FAT32?” The big deal is that HPFS was released well before FAT32, or even before Windows 9x was designed.

The disadvantage of HPFS is that it is not widely supported. Operating systems such as DOS and Windows cannot access HPFS volumes.

Managing Partitions and Volumes

For the A+ exam, you are required to know the steps to install a hard disk. After you have physically connected the drive and configured the jumper settings, you need to be aware of the steps to configure the partitions on the disk. The following is the order in which you configure the partitions on the disk:

- 1. Create a primary partition.**
- 2. Create an extended partition.**
- 3. Create a logical drive in the extended partition.**
- 4. Format the drives to create a file system.**

Understanding the order of the steps for partitioning and formatting a disk is important for the A+ Certification exam. Be sure to memorize the above list in order to correctly answer exam questions about preparing the hard disk.



A *partition* is defined as a segment of the hard disk, created by dividing the disk logically into discrete units. You create partitions for a number of reasons — you may partition a disk to organize your applications and operating system on drive C while storing your data on drive D. You may also partition a disk for more technical reasons, such as to run multiple operating systems on the same machine.

Whatever the reason for creating a partition, how you create and manage partitions is important for the A+ exam. This section examines different types of partitions and provides steps for creating, deleting, and formatting them.

Frequently, a partition is a means of providing better access to the information stored on a disk. For example, telling the kids that their games are on the D drive is usually easier than describing a complex path to the folder that holds the games.



You are limited to four partitions per disk, so be sure to plan them carefully.

Operating systems such as DOS, Windows 9x, and Windows 2000/XP/2003 (that are using basic disks — more on that in the section “Creating partitions and volumes in Windows 2000/XP/Server 2003,” later in this chapter) can create two types of partitions: primary partitions and extended partitions.

Primary partition

This is the partition that the computer boots from; the operating system’s boot files are loaded from here. You are allowed to have four primary partitions per disk. Because you may have multiple primary partitions (say, if you’re running several operating systems on the same computer), you must designate one primary partition as the *active partition* — the partition from which your normal operating system loads.

Extended partition

An extended partition allows you to extend beyond the four-partition barrier by being a partition that contains one or more *logical drives*. A logical drive is a block of disk space that is assigned a drive letter. As an example on how you could use extended partitions, you could set up three primary partitions and then decide that you would like to divide the last chunk of free space into three additional parts (for a total of six partitions). If you create another primary partition out of some of the free space, then you will have four parts — and that is your limit, four partitions per disk. What you can do instead is create an extended partition out of the remaining space after the three primary partitions have been created and then create three logical drives inside the extended partition. Logical drives are not partitions, so you are not limited to four. This will give you your six desired parts.

An *extended partition* is, in effect, the space that remains after the primary partitions are defined. The extended partition does not have an actual drive letter assigned to it; it's simply a container that holds all the logical drives that you build. A *logical drive* is a logical division of the hard disk that the computer treats as if it were a separate disk drive; it's the actual area of the extended partition to which documents are saved.

As an example, suppose you're partitioning a 6GB hard drive using the FAT file system. FAT cannot define partitions larger than 2GB, so you have to divide this drive into at least three different partitions: The first partition you define is the primary partition — a 2GB partition that also becomes the active partition (drive C). What's left is a 4GB extended partition that can store two logical drives (D and E), each of which can be no larger than 2GB. Figure 5-21 shows this partition configuration.

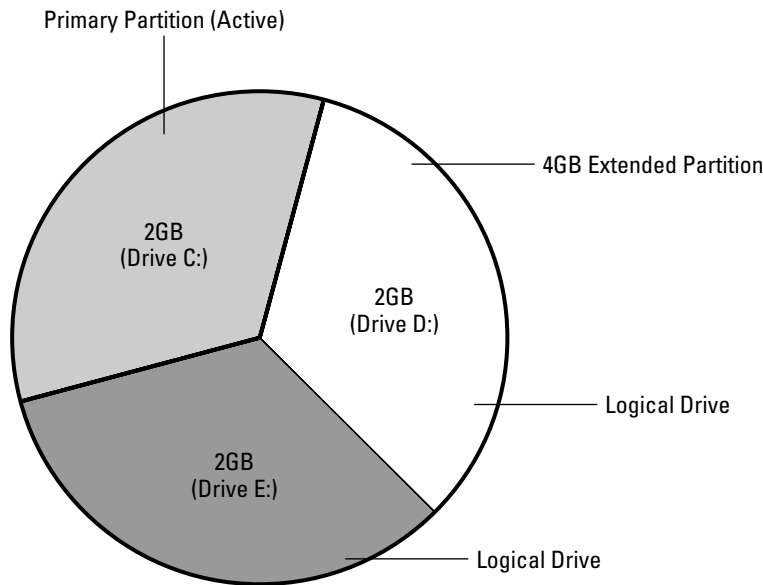


Figure 5-21: Partitioning a hard disk.

Note that the extended partition itself has no drive letter assigned to it. The extended partition is just a container to hold the logical drives — and they take the drive letters. Users of the system will be able to store data on drive C, drive D, or drive E!



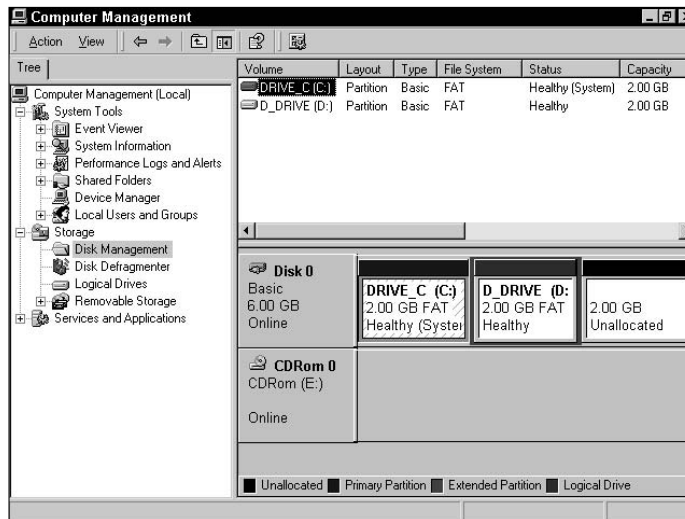
A hard disk can contain no more than four partitions, only one of which can be the extended partition. This means you could have three primary partitions and one extended partition to hold any logical drives. Having three primary partitions also shows why you have to set the active partition. A primary

partition is a bootable partition. But if I have three primary partitions, which one do I boot from? The answer is simple — the one defined as the *active partition*. Note that when you create the partition during the installation of Windows the partition is automatically marked as being the active partition.

Creating partitions and volumes in Windows 2000/XP/2003

To create partitions in Windows 2000/XP/Server 2003, you use a partitioning tool known as the *Disk Management snap-in*. You can open the Disk Management snap-in by choosing Start → Control Panel, clicking Performance and Maintenance, clicking Administrative Tools, and double-clicking Computer Management. Then, in the Computer Management window, select Disk Management. Figure 5-22 shows the Disk Management snap-in from Windows 2000/XP/Server 2003.

Figure 5-22:
Disk Management snap-in found within Windows 2000/XP/Server 2003.



When managing a disk in Windows 2000/XP/Server 2003, be aware that there are two types of disks: basic and dynamic.

Basic disks

Basic disk is the term Microsoft uses to describe a disk that supports partitions and all the limitations of partitions. If you can create partitions on a disk in Disk Management, then you are working with a basic disk. A basic disk has the following characteristics:

- ◆ **The disk is divided into partitions**
- ◆ **You are limited to four partitions per disk**
- ◆ **You are limited to one extended partition per disk**
- ◆ **You create primary partitions, extended partitions, and logical drives**
- ◆ **A basic disk is the default disk type in Windows 2000/XP/2003**

IT professionals have become accustomed to these characteristics when preparing disks for DOS, Windows 9x, and Windows NT. With Windows 2000 and above, you still have these limitations when working with a basic disk, but you may convert the basic disk into a dynamic disk. The benefit of converting to a dynamic disk is that you no longer work with partitions, so there are no partition limitations! You find out more about dynamic disks in the next section, but first you need to see how to create partitions on basic disks.



The following steps demonstrate how to create a partition on a basic disk within Windows XP. The steps are similar in Windows 2000 and Windows Server 2003.

- 1. Click Start, right-click My Computer, and choose Manage.**
- 2. When the Computer Management console has started, select Disk Management on the left side of the screen.**
On the right side, on the bottom half of the screen, notice that the disk type for Disk 1 is Basic Disk.
- 3. Right-click the unallocated space and choose New Partition, as shown in Figure 5-23.**

Figure 5-23: Creating a partition within Windows 2000/XP/2003.



The New Partition Wizard starts and displays a welcome screen.

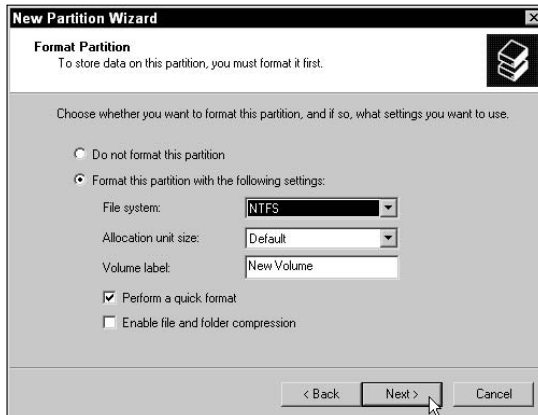
- 4. Click Next.**
The wizard asks what type of partition you will be creating, as shown in Figure 5-24.
- 5. Select Primary Partition and click Next.**

Figure 5-24: Choosing a partition type within the New Partition Wizard.



- 6. Type the desired size of the primary partition in megabytes and click Next.**
- 7. Choose the drive letter you want to associate with this partition and then click Next.**
- 8. Choose the file system you want to format the partition as (shown in Figure 5-25) and then click Next.**

Figure 5-25: Choosing the file system within the New Partition Wizard.



- 9. Click the Finish button.**

You will notice that the drive starts to format in the background, and it indicates the percent complete in the Disk Management utility.

- 10. When Disk Management is finished formatting the drive, close the Computer Management utility.**

Dynamic disks

A *dynamic disk* doesn't use partitions but rather *volumes* as discrete units of space. Because you are creating volumes instead of partitions, you don't have the four-partition limitation of basic disks. With dynamic disks, you are allowed to create as many volumes as you wish.

When creating a volume on a dynamic disk, you can create a number of different volume types:

- ◆ **Simple volume:** A simple volume is just a block of space that is similar in concept to a partition.
- ◆ **Striped volume:** A striped volume is a volume that is made up of equal space across multiple hard disks. With striped volumes, when you save a file to the volume, the file is saved across both disks at the same time. The benefit of a striped volume is a performance benefit from the fact that multiple disks are working at the same time to save the file.
- ◆ **Spanned volume:** A spanned volume is a volume that is made up of unequal amounts of space that span multiple disks. The benefit of spanned volumes is that you can join multiple areas of free space to create a single volume that users can access through a single drive letter.
- ◆ **Mirrored volume:** A mirrored volume is a volume that is made up of two disks. Data that is written to the volume is stored on both disks, each with a full copy of the data. If one of the disks should fail, the other disk has a copy of the data.

Mirrored volumes are only supported on the server versions of the Windows operating systems.

- ◆ **RAID 5 Volume:** A RAID 5 volume uses between 3 and 32 disks. Data saved to the volume is spread across all disks in the volume, along with parity data. The *parity data* is used to calculate data that is unreadable due to a failed disk in the volume.

RAID 5 volumes are also available only on server versions of the Windows operating system, so you won't be able to create one on Windows XP.

Creating a volume on a dynamic disk in Windows 2000/XP/2003 requires you to first upgrade the basic disk to a dynamic disk. After you upgrade the disk to a dynamic disk, you will notice that the `create partition` command has changed to a `create volume` command in the Disk Management console. To upgrade the basic disk to a dynamic disk, follow these steps:

1. Click Start, right-click My Computer, and choose Manage.
2. When the Computer Management console has started, select Disk Management on the left side of the screen.



On the right side, in the bottom half of the screen, notice that the disk type for Disk 1 is Basic Disk.

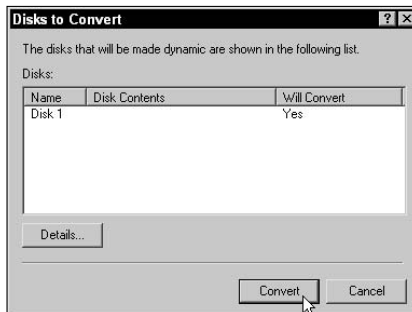
3. Right-click the disk and choose **Convert to Dynamic Disk**, as shown in Figure 5-26.

Figure 5-26:
Converting a basic disk to a dynamic disk.



4. In the **Convert to Dynamic Disk** dialog box, ensure that the disk number you wish to convert is selected and click **OK**.
5. Click the **Convert** button on the **Disks to Convert** screen, as shown in Figure 5-27.

Figure 5-27:
Confirming the disk to convert to a dynamic disk.



You get a warning letting you know that you will be unable to start older operating systems from the disk after it is converted.

6. Click **Yes** in the **Warning** dialog box.
7. Click **Yes** to **dismount any file systems being converted**.

After the drive is converted, you will notice that the legend has changed to include a simple volume. You will also notice that the disk is now a dynamic disk.

Now that you have converted a basic disk to a dynamic disk, you are ready to create a volume in Windows XP. Creating a volume is similar to creating a partition. To create a volume, follow these steps:

- 1. In the Disk Management console, right-click Unallocated Space and choose New Volume, as shown in Figure 5-28.**

Figure 5-28: Creating a new volume on a dynamic disk.

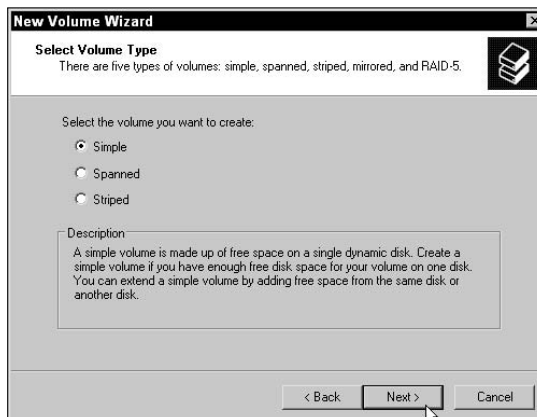


The New Volume Wizard appears.

- 2. Click Next.**

You are shown a list of volume types you can create (see Figure 5-29).

Figure 5-29: Selecting a volume type when creating a new volume.



- 3. Choose the volume type you wish to create and then click Next.**
- 4. Type in the amount of space in megabytes to use for the volume and then click Next.**
- 5. Select a drive letter for the volume and click Next.**
- 6. Specify a file system to use and choose to perform a quick format. Then click Next.**
- 7. Click Finish to create the volume.**

Formatting partitions and volumes

After you have created the partitions or volumes, your next step is to format these partitions or volumes so that you may start storing data on them. When you format the partitions or volumes, which now show as drive letters in the My Computer icon, you choose which file system to format them with. Before you format the partitions, you should review the different types of file systems and the advantages and disadvantages of each (see the section “Understanding File Systems,” earlier in this chapter).

Formatting a drive prepares the drive for storing information. The `format` command creates a root directory on the disk as well as two tables used to store information about the files and to aid in the retrieval of these files.

The first table is the *Directory Entry Table (DET)*; it lists all the files stored on the drive, along with the date that each file was last modified. It also stores the starting cluster for each file stored on the drive. The DET is used when the system goes to open a file; the system looks in the DET for the file. Once an entry is found for the file being opened the starting cluster for the file is determined and then the system goes to that cluster to retrieve the file contents.

The second table is the *File Allocation Table (FAT)*. The FAT lists each cluster, showing you which clusters are used and which clusters are free. The FAT also indicates any clusters that have been marked as bad clusters, which are unusable. When a file spans multiple clusters, the FAT shows that the first cluster is linked to the next cluster by indicating the next cluster value as part of the FAT entry. The last cluster that is used to store the data for the file is marked with an *end of file (EOF)* marker — this is how the system knows it has reached the last cluster for the file. Figure 5-30 shows a formatted drive along with the DET and FAT.

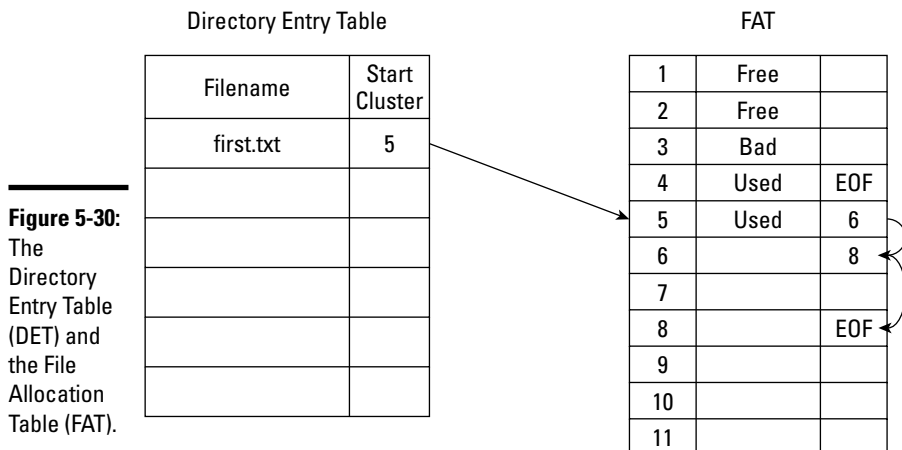


Figure 5-30: The Directory Entry Table (DET) and the File Allocation Table (FAT).


In Figure 5-30 a file named `first.txt` is listed in the DET with the beginning cluster location of 5. When a user requests the file, the operating system looks for the file in the DET. In this example, the DET states that the first cluster containing data for the file is cluster number 5. Then the system looks at cluster 5 in the FAT and sees that cluster 5 is being used and that it continues on to cluster 6 and on to cluster 8. The system knows that cluster 8 is the last cluster containing data for the file because of the EOF marker.

To format a partition or volume in Windows 2000/XP/Server 2003, open My Computer, right-click the drive letter that is assigned to the partition or volume, and choose Format.

With Windows 2000/XP/2003, you may also format the partition or volume from within the Disk Management snap-in, which is where you created the partition or volume to begin with. To format a partition/volume in Windows 2000/XP/Server 2003 using Disk Management, perform the following steps:

1. Click Start, right-click My Computer, and choose Manage.
2. When the Computer Management console has started, select Disk Management in the left side of the screen.
3. Over on the right side, in the bottom half of the screen, right-click the partition you want to format and choose Format, as shown in Figure 5-31.

Figure 5-31: Choosing the Format command in Disk Management.




When the Format dialog box appears, you may choose from a number of settings (shown in Figure 5-32) on how you wish to perform the format. The settings include

- *Volume Label:* Use this option to specify a label for the partition. For example, I usually have a drive that is labeled “Data” where I store all my data. This label is just a friendly name used to identify what you might be using the drive for.
- *File System:* Use this option to specify whether you want to use FAT, FAT32, or NTFS as the file system.

- *Allocation Unit Size*: Use this option to select the cluster size for the partition or drive.
- *Perform a Quick Format*: A *quick format* does not perform a surface scan on the disk to check for errors.
- *Enable File and Folder Compression*: You may implement compression on the drive by selecting this option. *Compression* allows you to save disk space by using less space on the disk for the data that is stored there. The compression feature shrinks the file when the file is saved and decompresses it when the file is opened.

Figure 5-32: Format options available in Windows XP when formatting a drive.



4. Select and/or enter your options and click OK to format the drive.

5. Click OK to format the drive.



TIP

You may also use the command-line version of the `format` command by going to a command prompt and typing `format`. For example, to perform a quick format of drive D and then apply a volume label of “Pictures,” the command is:

```
format d: /q /v:Pictures
```



Labs 5-3, 5-4, and 5-5 give you the opportunity to practice creating partitions and volumes in Windows XP. You can find these labs in the `Labs.pdf` file in the Author directory of the CD-ROM.

Securing Data with RAID

Windows servers have a built-in software implementation of *RAID* (*Redundant Array of Inexpensive Disks*) that you can take advantage of. RAID is a method of implementing redundancy — duplicated information — on your hard drives; if one disk fails, the other disk(s) can provide the missing information. RAID is a method to implement fault tolerance. *Fault tolerance* is the idea that if there is a hardware failure in the system, such as if a hard drive fails, the

system can continue to operate as normal. There are many different levels of RAID, but the only RAID levels that provide redundancy in Windows Servers are RAID level 1 and RAID level 5.



Windows Servers also support RAID level 0, called *disk striping* or *striped volumes*. No redundant information is stored on striped volumes, which is why striped volumes are classified as RAID level 0. (Think of it as a zero level of redundancy.)

Mirroring/duplexing (RAID level 1)

The type of hard drives that are normally found in servers are SCSI drives, which means that there is a SCSI adapter (controller) that connects the drives to the systems. *Disk mirroring* is the use of two disks on a single controller to create *full redundancy* — whatever is placed on one disk is copied to the second disk. When creating a disk mirror, you must use two disks so that if one disk fails you can rely on the copy of the data that is stored on the other disk.

When the mirror is established, you will have a new drive letter that's accessible from the My Computer icon. This drive actually shows the data stored on *both* disks; if you save a file to this drive, it's written to both disks that make up the mirror; you don't see two representations within the My Computer icon. But if you use the Disk Management console, you see that the two disks are part of the mirrored volume.

Disk duplexing is the same idea as disk mirroring but requires the installation of an additional SCSI controller. Disk mirroring is fault-tolerant if a drive fails (because the other drive is available), but there is no fault tolerance if the controller fails (because there is only one). If you add an additional controller and place one drive on one controller and the other drive on the other controller, you have a more fault-tolerant solution. If one drive fails, you have the other; if a controller fails, you have the other drive running off the other controller. Figure 5-33 shows the difference between a mirror and a duplex.

To create a mirrored volume on Windows Server 2003, you first need to ensure that you have converted both of the disks that you want to use as mirrors to dynamic disks. See the previous section, "Dynamic disks," to find out how to do this.

After you have converted both of the disks to dynamic disks, you can create a mirrored volume by following these steps:

- 1. Click Start, right-click My Computer, and choose Manage.**
- 2. In the Disk Management console, right-click the unallocated space and choose New Volume (as shown in Figure 5-34).**

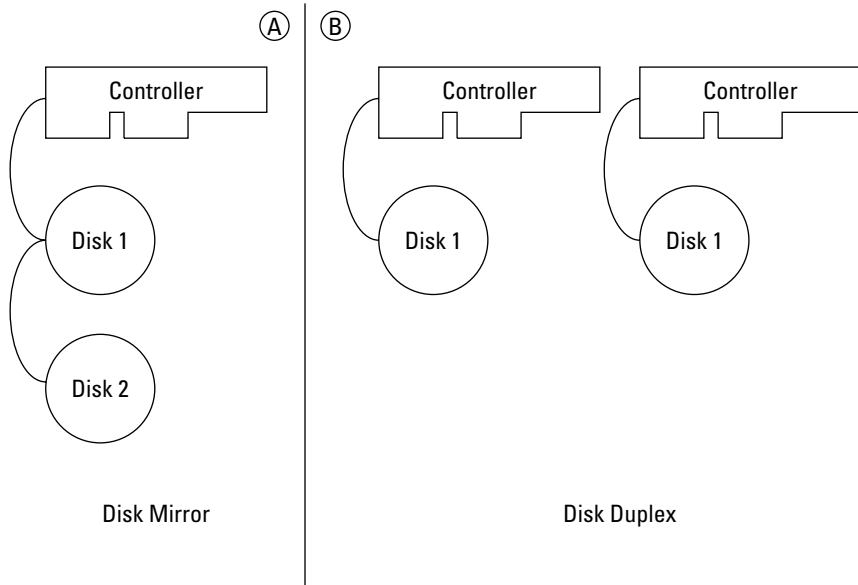


Figure 5-33: Comparing disk mirroring with disk duplexing.

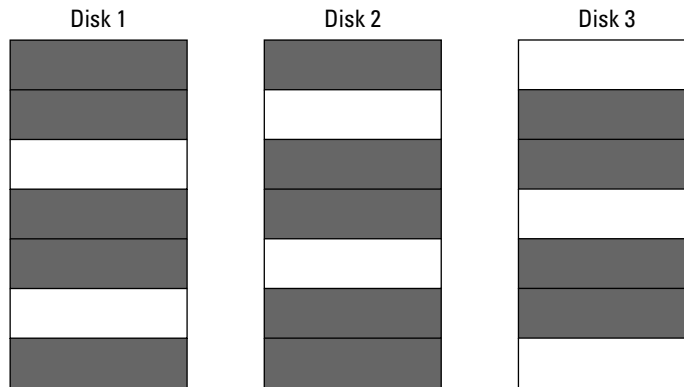


Figure 5-34: Creating a new volume in disk management.

Legend:
 ■ Data
 □ Parity Information

3. When the welcome screen to the New Volume Wizard appears, click Next.

You are asked what type of volume you wish to create.

4. Choose Mirrored and then click Next (as shown in Figure 5-35).

Figure 5-35: Selecting a mirrored volume type when creating a new volume.

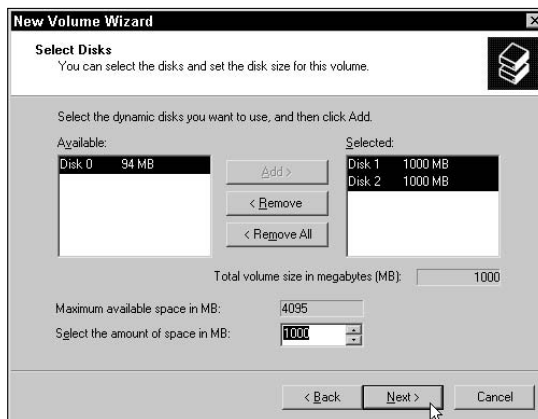


You are next asked to select which disks will be members of the volume and how much space you wish to use on each disk.

- 5. From the available disks on the left side, choose the disk you wish to have the mirrored content stored on and then click Add.**

Notice that the original disk that you right-clicked to create the volume is already selected on the right side. In Figure 5-36, I have selected Disk 1 and Disk 2 to participate in the Mirrored Volume.

Figure 5-36: Selecting the disks for the mirrored volume and setting the volume size.



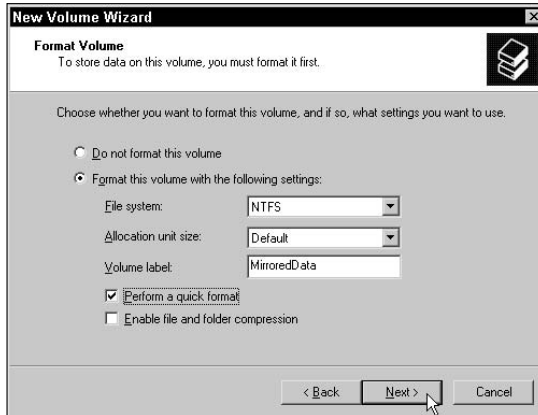
- 6. Specify how much space to use for the volume in the Select the Amount of Space in MB text box and click Next.**

In Figure 5-36, I use 1000MB.

- 7. Choose a drive letter for the volume and then click Next.**

8. Select NTFS as the file system and type a volume label of `MirroredData`. Also select to perform a quick format (as shown in Figure 5-37) and then click Next.

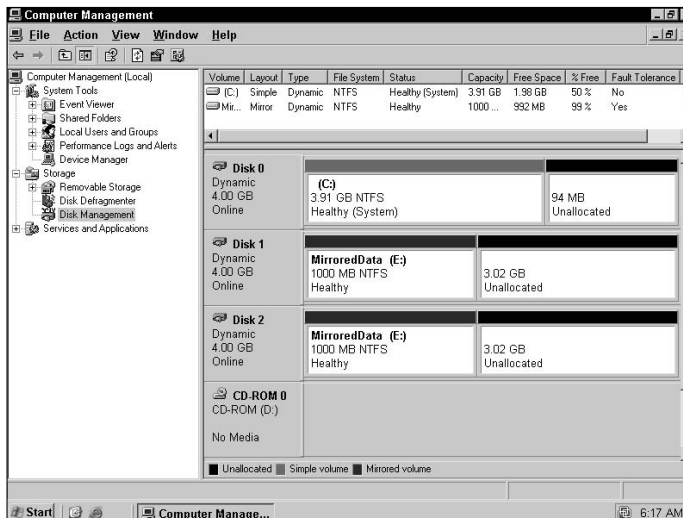
Figure 5-37: Assigning a volume label to the mirrored volume.



9. Click Finish.

You have now created a mirrored volume, so anything that is stored on the drive will get stored on both disks. If one disk fails, you have a copy of the data on the other disk. Notice in Figure 5-38 that the mirrored volume takes the same drive letter (E: in this case) for both disks. Therefore, if you store something on that drive, it is written to both disks.

Figure 5-38: Viewing the mirrored volume in disk management.



RAID 5 Volume (RAID level 5)

The problem with RAID level 1 is that you essentially waste half of the money you spend on hard drives because, under normal conditions, you only use one disk. A RAID solution that you can use where you get more disk space for your dollar is known as *RAID level 5*, or a *RAID 5 volume*.

Microsoft's implementation of RAID 5 volume uses a minimum of 3 disks and can use up to 32 disks. When data is saved to the RAID 5 volume, the information is written across all disks in the array.

A RAID 5 volume also stores parity information on a different disk for each write operation. This parity information is an “answer” that is generated after the data being written is run through an algorithm. The answer is then stored on one of the disks for that write operation. If a disk fails and a piece of data cannot be retrieved, the data is recalculated based on the answer (parity data) that is stored in the array.

Sound confusing? Here's a simple example: What is the value of x in the following formula?

$$4 + x = 9$$

You probably had no problem coming up the answer of $x = 5$. You simply subtracted the known data from the given answer to calculate what is missing. RAID 5 volumes do the same thing with the parity data (the “answer”) when there is a disk failure. When one disk fails, the fault-tolerant disk driver simply subtracts the known data (from the existing disks) from the parity data to generate the missing data on the fly.

A popular question I get is this: “Will RAID 5 volumes be able to recover from multiple disk failures?” To answer that, look at a simple algebra equation. What is the value of x in the following formula?

$$y + x = 9$$

We cannot say for sure what the value of x is because there are now two variables involved. The software implementation of RAID that is built into the operating system can calculate the missing data only when there is one failed drive.

Now take a look at how to create a RAID 5 volume in Windows Server 2003. Remember that a RAID 5 volume must use at least three disks — so be sure to have three disks with unallocated space. Then follow these directions:

- 1. In the Disk Management console, right-click unallocated space and choose New Volume.**

2. When the welcome screen to the New Volume Wizard appears, click Next.

You are asked what type of volume you wish to create.

3. Choose RAID 5 (as shown in Figure 5-39) and then click Next.

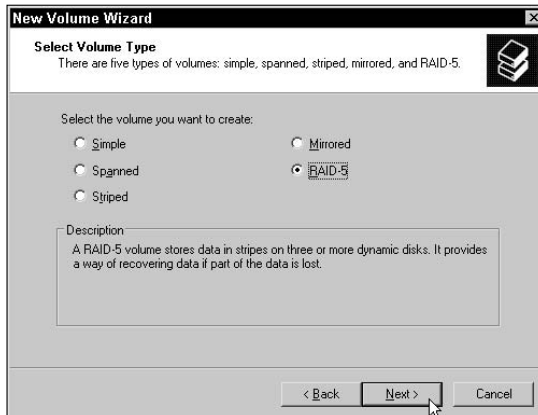


Figure 5-39:
Choosing to
create a
RAID 5
volume.

4. Choose at least two additional disks to be members of the RAID 5 volume by highlighting each disk on the left side of the screen and clicking Add.

In my example, I chose Disk 0, Disk 1, and Disk 2.

5. Specify how much space to use on each disk in the Select the Amount of Space in MB text box.

Note that I used 90MB on each disk, and my total space for the volume is 180MB (see Figure 5-40). But wait a second . . . that doesn't make sense. There are three disks in the RAID 5 volume, and if each disk uses 90MB, then the volume size should be 270MB, shouldn't it?! The answer is no, the volume size is only 180MB because one-third of the space is used for parity data.

6. Click Next.

7. Choose a drive letter for the volume and then click Next.

8. Select NTFS as the file system and type a volume label of RAID5Data. Select to perform a quick format and then click Next.

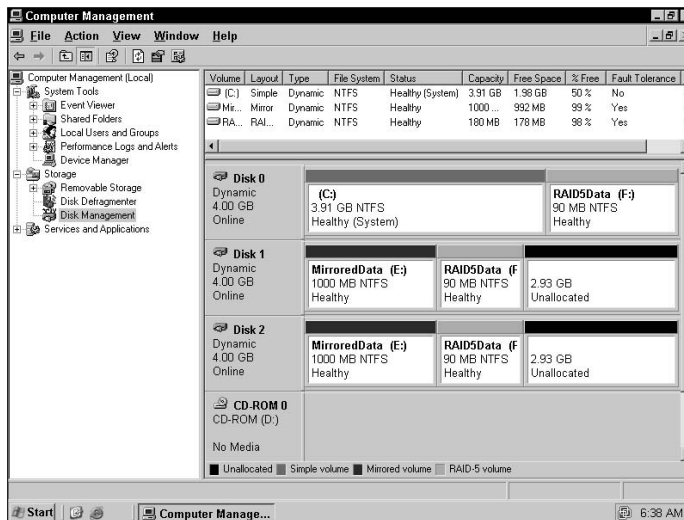
9. Click Finish.

Figure 5-41 displays the volume as it appears in Disk Management.

Figure 5-40: Select the disks to be members of the RAID 5 volume and specify the size of the volume.



Figure 5-41: Viewing the RAID 5 volume in disk management.



Understanding Management Tools

After partitioning the disk and formatting the drive, some maintenance still needs to be done on a regular basis. Drive maintenance helps you address two areas of concern that have an impact on a user's data:

- ◆ **File system optimization:** This means getting the best performance from your file system. If you can open the files stored on a computer but they open much more slowly than they did two months ago, your system needs optimizing.

- ◆ **Data integrity:** This means having files that are uncorrupted so the data is intact and accessible. Sometimes a file cannot be opened because the file system has lost parts of the file or portions of the disk are bad — in which case you are not concerned with performance but instead concerned about being able to open the file.

The following sections describe some operating system tools that can help you accomplish these two goals as you maintain your system.

Defragmentation utility

After a drive is formatted and you start storing information on it, the information is written to one cluster at a time. This means that on a freshly formatted drive, the contents of a file are written to clusters on the disk that are side by side. This ensures optimal performance when opening a file because the read/write heads don't have to jump from one end of the disk to another to open a single file. Unfortunately, the disk won't stay in this state because as you add to and delete files, the contents of these files are scattered throughout the disk. The following minitable shows a fragmented disk.

1	2	3	4
5	6	7	8

In this example, you can see that block 2, 5, and 7 belong to the same file but are scattered throughout the disk. This causes a performance decrease when accessing the file because the read/write heads need to locate the contents of the file that are spread throughout the disk. Disk defragmenting applications clean this up by taking all the data of a single file and placing it in clusters that reside side by side. You can see below what the disk would look like after a Windows defragmentation.

1	2	3	4
	2	3	4
5	6	7	8

Windows 2000/XP/2003 has a Disk Defragmenter utility in the Computer Management console. This great little tool has an analyzer that you can run first — it checks the selected disk and reports where the used space is and where the free space is. This tells you whether you need to defragment. Figure 5-42 shows the Disk Defragmenter utility in Windows XP after analyzing drive c.

After the hard drive has been analyzed, you can also view a report that shows you detailed information about the analysis just performed. When you're happy with the analysis, you can perform the defrag by clicking the Defragment button. Figure 5-43 shows the report generated by the analysis of drive c.

Figure 5-42: Defragmenting a disk in Windows 2000/XP.

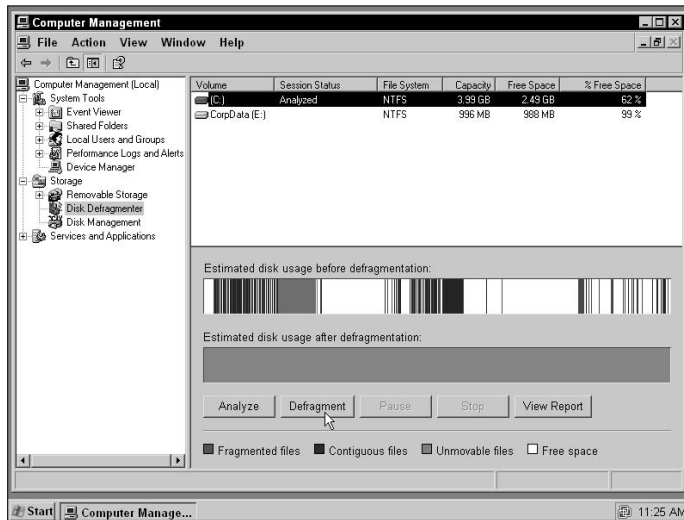
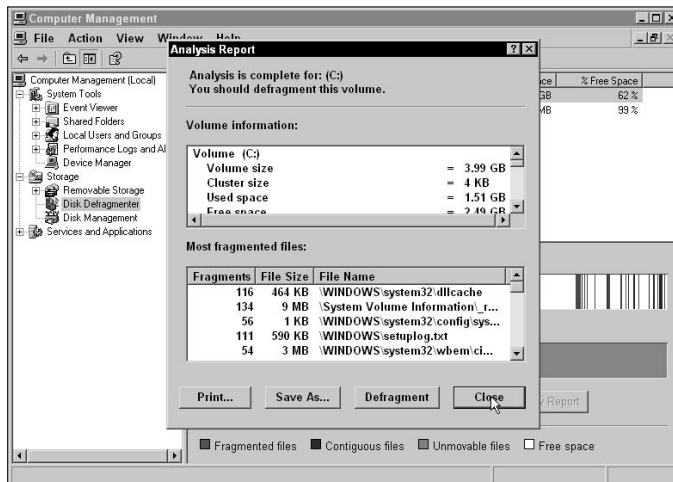


Figure 5-43: The analyzer's report in Windows 2000/XP/2003.



Check Disk utility

When a drive is formatted, two FAT tables are created — the Directory Entry Table and the FAT table. The FAT table links all the clusters that make up a file. What if that link is lost, or points to the wrong location? If the file system loses the link that joins two clusters together, the file becomes unreadable and has lost its data integrity. *ScanDisk* is a Windows utility in older

Microsoft operating systems that has been replaced by the Check Disk option in Windows 2000, XP, and Server 2003. The Check Disk option not only looks for lost links but also scans the disk surface for bad blocks (clusters that cannot be written to), marking them as bad so they are not used. To perform an integrity check on a drive in Windows XP, follow these steps:

1. Open the My Computer icon and select the drive you want to check.
2. Right-click the drive and choose Properties.
3. On the Tools tab, click the Check Now button.

Figure 5-44 shows the dialog box that appears.

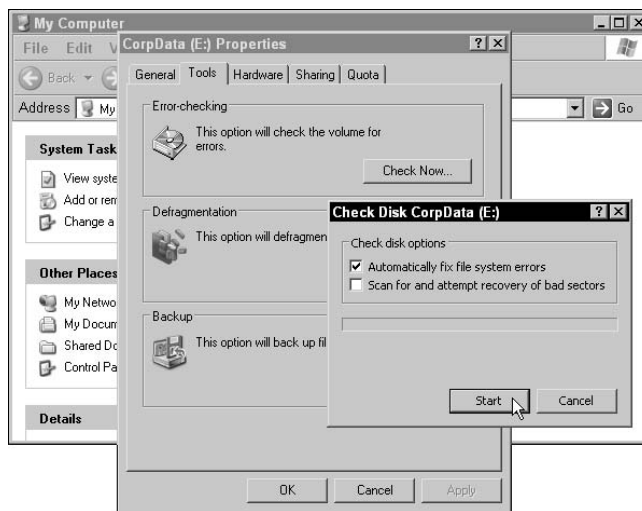


Figure 5-44:
The Check
Now option
in Windows
2000/XP/
2003.

You may choose to automatically fix any file system errors and specify whether you want Windows XP to scan the surface of the disk and attempt to fix any problems with sectors.

4. Choose the options you want and then click Start.

The scan may take a few minutes and when completed will display summary information.

Disk Cleanup


Another important tool for Windows XP is the Disk Cleanup tool. The Disk Cleanup tool scans a disk for files that can be safely removed from your system in order to free up disk space. The Disk Cleanup utility can remove a

number of different types of files to help free up disk space. Disk Cleanup can remove temporary Internet files, Windows temporary files, and applications no longer used.

To perform a disk cleanup, follow these steps:

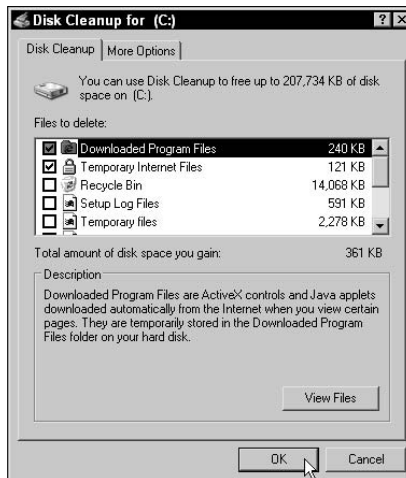
1. Choose **Start**→**All Programs**→**Accessories**→**System Tools**→**Disk Cleanup**.
2. Select the drive you wish to clean up, as shown in Figure 5-45.

Figure 5-45: Selecting the disk to clean up.



3. Click **OK**.
4. Select which files should be deleted to reclaim disk space, as shown in Figure 5-46.

Figure 5-46: Choose which types of files to remove with Disk Cleanup.



5. Click **OK** and then **Yes** when asked if you are sure.

Unneeded files are then removed from your system to reclaim the space.

Getting an A+

This chapter illustrates the importance of hard disk management and the utilities used to perform that management. The following are some key points to remember when managing hard disks:

- ◆ A *cluster* is the allocation unit for a file.
- ◆ The cluster size for a partition is based on the file system being used and the size of the partition.
- ◆ The FAT file system is limited to a maximum partition size of 2GB. FAT32 increased the maximum allowable partition size to 2000GB.
- ◆ Windows NT implemented its own file system, called NTFS. NTFS has a number of benefits, some of which are security, auditing, and compression. Windows 2000/XP/2003 use NTFS 5.0, where the Encrypting File System and quotas have been added to the list of reasons you should be using NTFS.
- ◆ To optimize your drive, run the defragmentation tool often.
- ◆ To verify the integrity of the drive, run ScanDisk (Window 9x) or Check Now in Windows 2000/XP/Server 2003 often.
- ◆ A primary partition is the bootable partition for the system and must be set active. An extended partition holds logical drives for storing information.
- ◆ You can remove temporary files with the Disk Cleanup utility in Windows XP.
- ◆ SCSI devices need to have a host adapter and unique IDs, and the SCSI bus must be terminated at both ends.
- ◆ SATA drives are much faster than IDE and support hot swapping.
- ◆ RAID 1 is disk mirroring/duplexing.
- ◆ RAID 5 stores the data along with parity information.

Prep Test

- 1** What is the size of a sector on a hard disk?
 - A 512 bytes
 - B 1K
 - C 4K
 - D 512K

- 2** How many devices can exist in an IDE chain?
 - A 1
 - B 2
 - C 3
 - D 4

- 3** What physical component of the disk is responsible for reading and writing data on the disk?
 - A Platter
 - B Sector
 - C Cluster
 - D Read/write head

- 4** How many wires can be found in an IDE ribbon cable for a hard disk?
 - A 34
 - B 40
 - C 50
 - D 66

- 5** In a SCSI chain, what devices are required to be terminated?
 - A The first device in the chain
 - B The last device in the chain
 - C Devices at either end of the SCSI chain
 - D Any hard disk in the SCSI chain

- 6** In a SCSI chain, what device should take a SCSI ID of 0?
 - A CD-ROM
 - B The bootable (first) hard disk
 - C The SCSI host adapter
 - D Printer

- 7** In an IDE chain, CD-ROM devices are usually what?
- A Masters
 - B Terminated devices
 - C Devices assigned ID 0
 - D Slaves
- 8** Which of the following describes the cabling for internal SCSI?
- A 40-wire ribbon cable
 - B 34-wire ribbon cable
 - C 50-wire ribbon cable
 - D None of the above
- 9** Which of the following best describes where the Master Boot Record is located?
- A Last track on the disk
 - B First track on the disk
 - C First sector on the first track of the first side of the first platter
 - D First sector on the last side of the last platter
- 10** How many devices can exist in an EIDE chain?
- A 1
 - B 2
 - C 3
 - D 4
- 11** Which of the following tools will help optimize the file system?
- A Backup
 - B Defragment
 - C ScanDisk
 - D Format
- 12** What is the maximum capacity for a partition using the FAT file system?
- A 2MB
 - B 2GB
 - C 16GB
 - D 2000GB
- 13** What type of partition is used to boot the computer?
- A Extended
 - B Primary
 - C Mirror
 - D Stripe set with parity

- 14** What is the allocation unit for the storage of a file?
- A Sector
 - B Track
 - C Partition
 - D Cluster
- 15** Which file system includes the encrypting file system?
- A HPFS
 - B FAT
 - C NTFS 5.0
 - D FAT32
- 16** What utility manages partitions in Windows XP?
- A Disk Administrator
 - B Disk Management
 - C Backup
 - D fdisk
- 17** What type of partition must be flagged as being active?
- A Extended
 - B Logical
 - C Primary
 - D All partitions
- 18** What types of volumes are available on a dynamic disk? (Select all that apply.)
- A Primary
 - B Simple
 - C Extended
 - D Spanned
- 19** You don't seem to be able to create a simple volume on Windows XP; why not?
- A You are not using NTFS as the file system.
 - B You need to format the drive.
 - C You must have an extended partition before you can create a simple volume.
 - D You need to convert the basic disk to a dynamic disk.

Answers

- 1 A. A single sector on the disk is 512 bytes. A group of these sectors makes up a cluster, which typically could be 4K, 16K, or 32K. *See "Sectors."*
- 2 B. IDE supports only two devices in an IDE chain. EIDE improves on this limit by allowing four devices. *Review "IDE/ATA standard."*
- 3 D. The read/write heads are responsible for moving over the disk surface and writing data to the disk. A platter is one of the disk plates, a sector is where the data is written, and a cluster is a group of sectors. *Check out "Read/write process."*
- 4 B. The hard disk is connected to the motherboard with a 40-wire ribbon cable. A 34-wire ribbon cable is used to connect floppy drives, and a 50-wire ribbon cable is used for internal SCSI devices. *Peruse "IDE cabling."*
- 5 C. Devices at each end of the SCSI bus must be terminated. This is usually the SCSI host adapter plus the last device in the chain. *Take a look at "Termination."*
- 6 B. A bootable hard disk should take SCSI ID 0, while the host adapter usually takes an ID of 7. The other devices can be assigned any number in the chain as long as it is a unique number. *Peek at "Addressing."*
- 7 D. IDE CD-ROMs are slaves in an IDE chain. IDE devices do not get assigned IDs and do not have to be terminated. *Look over "Master/slave configuration."*
- 8 C. An internal SCSI connector uses a 50-wire ribbon cable, an IDE device uses a 40-wire ribbon cable, and a floppy drive uses a 34-wire ribbon cable. *Study "SCSI cabling."*
- 9 C. The master boot record is the first sector on the first track of the first side of the first platter. *Refer to "Sectors."*
- 10 D. An EIDE chain can have up to four devices. IDE only supports two devices, and the original SCSI supports eight devices in the chain. *Examine "EIDE/ATA-2 standard."*
- 11 B. A defragmentation will optimize the file system. It will reorganize the clusters on the disk so that all the data that makes up a file is in the same area so the read/write heads will not have to jump around from one end of the disk to the other. *See "Defragmentation utility."*
- 12 B. FAT file systems can only access 2GB partitions. This has seen a big improvement with FAT32, which can access 2000GB partitions. *Review "The FAT file system."*
- 13 B. A primary partition is the type of partition that is bootable. Extended partitions are used to hold logical drives. Stripe set with parity is used as a RAID solution on servers. *Check out "Managing Partitions and Volumes."*

- 14** D. Files are written to clusters on the disk. Only one file can occupy a cluster, regardless of whether the file fills the entire cluster. *Peruse "Clusters."*
- 15** C. Windows 2000's version of NTFS, called NTFS 5.0, implemented the Encrypting File System. *Take a look at "NTFS 5.0."*
- 16** B. The Disk Management utility in Windows 2000 allows you to manage partitions. *Peek at "Managing Partitions and Volumes."*
- 17** C. Primary partitions are the type of partitions that are flagged as being active. If you have two or three primary partitions on your drive, the primary partition that is flagged as being active will be the partition the system boots from. *Look over "Managing Partitions and Volumes."*
- 18** B, D. Of the choices provided, Simple and Spanned are the only volume types available on a dynamic disk. Primary and Extended are partition types and are only available on basic disks. *Study "Dynamic disks."*
- 19** D. In order to create a volume, you must ensure that you have converted the basic disk to a dynamic disk. After you have converted to a dynamic disk, you can create any of the different types of volumes. *Refer to "Dynamic disks."*

Chapter 6: Working with Power

Exam Objectives

- ✓ Understanding the features, functions, and installation of ATX power supplies
- ✓ Identifying features and functions of AC adapters

This chapter takes a look at two main items, *Power Supply Units (PSU)* and *Uninterruptible Power Supplies (UPS)*. Power supplies should actually be called *power converters* because the largest single function they perform is a conversion of power from 120V or 240V to something that can be used by the computer system.

As a CompTIA A+ Certified Professional, you will at least have to add a component to a computer that requires power from the power supply, such as a hard drive, and you will also probably be required to troubleshoot and replace a power supply in the computer system. The information in this chapter gives you all of the technical specifications and connector types that you will need to know about a power supply, and together with the troubleshooting chapter (Book IV, Chapter 2), you will have all you need to know about power supplies.

Knowing the Basics of Power Terminology

There are four terms that you should learn in order to deal with power in computer systems. These four terms are *volts*, *amps*, *ohms*, and *watts*, and, if you have a firm grasp on how they relate to each other, you will be able to understand how power is measured and used within computer systems. Electrical current flows with the same principles that are used when water flows, so water is often used as an analogy for electrical current.

Volt

A technical definition of a *volt* is a unit of electrical potential difference or the potential difference across a conductor when a current of one ampere dissipates one watt of power. If we put this into water terminology, then voltage is represented by the pressure that is the water supply tank. This is a potential because, unless you provide an exit path for water to leave the tank, this potential is not realized. If you increase the pressure in the tank, then you increase the potential to supply water or power.

Ohm

An *ohm* is a measure of resistance or electrical impedance. When comparing electricity to water, resistance is determined by the size or diameter of the hose or pipe that is used to supply water, where a smaller hose produces more resistance.

Amp

An *ampere* or *amp* is a unit of electric current or the measure of electrical flow. When referring back to water, it is also the rate of water flow. There is a direct relationship between volts, resistance, and current in that, if you reduce resistance or increase voltage, you will increase the current. This relationship is called Ohm's law, and is represented by this formula:

$$I = V / R$$

In this formula, I is current (Amps), V is voltage (Volts), and R is resistance (Ohms).

Watt

The technical definition of a watt is that it is a unit of power, equal to one joule per second. If we equate this to water again, then watts would have you use the water that is flowing to do something. If this water were to run over a waterwheel, then watts would be the measure of how fast the wheel is able to be turned by the water. If you increase the current, then the wheel is able to turn faster. The formula to calculate watts is:

$$P = V * I$$

In this formula, P is power (Watts), V is voltage (Volts), and I is current (Amps).

In general, when you are working with computer systems, you will be concerned with the voltage supplied by the outlet that you plug your equipment into, the current that your equipment draws from the outlet, and the power that is used by devices that are connected to a computer power supply.

Identifying the Purpose of Power Supply Units



Remember the tasks and specifications listed in this section.

A computer's power supply is responsible for several tasks. The following list contains typical tasks and specifications related to power supplies:

- ◆ *Conversion* of voltage from the building-supplied AC voltage to various DC voltages used by components inside the computer.
- ◆ *DC Voltage Regulation* to provide a very small tolerance of error to the devices receiving the power. Typically a +/-5% voltage differential is allowed through the specification for 12VDC, 5VDC, and 3.3VDC connections.
- ◆ *Over Current Protection* specifies that the power supply should handle some level of abnormally high current to prevent fuses in the power supply from blowing unnecessarily.
- ◆ *Input Under Voltage* specifies that if the voltage being supplied by the building drops below the normal operating level, it should not cause damage to the power supply itself — it will, however, likely cause your computer to turn off. For a 120V power supply, this low voltage level is usually 100V, so at voltages below that level, the power supply will simply turn off.
- ◆ *Energy Star* is a specification that defines how much power the power supply should provide to the computer components during reduced power states, such as when the computer is put to sleep or on standby.
- ◆ The *PS_ON#* feature allows the motherboard to control the power supply through a soft power switch. Power supplies for ATX motherboards do not use a physical switch like your light switch, which allows a circuit to be open or closed; but rather your computer's power button is connected to the motherboard, and when you press the button, the motherboard determines what action will be taken. The power supply allows the motherboard to control the power state through the *PS_ON#* feature.
- ◆ Depending on the form factor of the power supply, the specification for the power supply may either recommend or require the use of cooling fans.
- ◆ Input supply voltage (120V or 240V) may be static or configured via manual switching or automatic switching. Many laptops have automatic switching, while most desktop power supplies have manual switching.

If you do not see a manual switch on your power supply, then you will want to read the input voltage range that will be displayed on a label on the power supply. For desktop power supplies, this label will usually be found in an inconvenient location on the power supply, and you will likely need to remove the power supply from the case to read the label. In the case of laptop power supplies, even though it may have a switch, you must read the label to ensure that it does support both 120V and 240V.

Automatic switching just means that the power supply will detect and use the voltage setting for the power that is supplied. Remember back to the start of this chapter, if you increase the voltage, then you increase the power; so the power supply not only increases the voltage, but also reduces the current inside the power supply, to allow the internal devices to have a consistent power level.

Identifying Power Supplies

The most common type of motherboard that you will have to provide power to will be some form factor of the ATX family (see Book II, Chapter 1), and, because of that, some form of power supply that is compatible with an ATX motherboard will be required. In the following sections, I spend most of the time discussing the things that are common to all power supplies that are compatible with ATX motherboards, but I also take a look at the different form factors that exist for power supplies so that the many sizes and shapes of computer cases may be supported — from tower, to small form factor (SFF), to slim line cases, and all of the sizes in between.



When the name ATX power supply is used properly, it refers to a power supply that provides power to an ATX motherboard, and it has very specific size characteristics. There are many other power supplies, such as CFX, LFX, SFX, and TFX, that have different size characteristics, but all provide power to the ATX family of motherboards. Each size and shape of power supply has its own specification. In this chapter, I will only talk about power supplies that provide power to ATX motherboards, regardless of their size and shape. All of these power supplies will use the ATX-compatible power main power connector, which makes them different from the older AT power supply that was used with older generations of motherboards.



If you want to find out more about the specifications for different power supplies that provide power to ATX motherboards, look at Intel's Power Supply Design Guide for Desktop Platform Form Factors document at www.formfactors.org. This Intel-operated site has specifications for many of the standard form factors used for computers and components.

The ATX power supply has undergone some basic changes in power that is supplied and in the format of the connectors. The current version of the ATX specification is 2.2. The biggest difference between the version 1.3 ATX specification and the 2.x ATX specifications is the use of a 24-pin main power connector used in the 2.x ATX specification, rather than a 20-pin connector under the 1.3 ATX specification.

With ATX power supplies having the same types of power connectors, one big difference between the power supplies you could purchase is how much power they provide to connected devices. The power supply rating refers to the total wattage that the power supply can provide to devices within your computer. While ATX specification allows for power supplies to have ratings from 250 watts up to large ones at 800+ watts, most are typically around 500 watts to 600 watts. Lower wattages are available, too, but these usually follow the SFX specifications (sometimes called Micro ATX power supplies), and these have power output down to 160 watts.

The power supply rating required for any computer depends on the devices within the computer that need power. In most cases, bigger is better because systems that don't have enough power for devices can experience strange, intermittent problems (see Book IV, Chapter 2). Table 6-1 provides a listing of some basic systems devices and power usage estimates as a guide to sizing up a power supply, but remember to always leave yourself some extra room.

Device	Power
AGP video card	30–75W
PCI video card	30–35W
AMD Athlon XP 1.5MHz–2.5GHz	66–77W
AMD Athlon 64 3.0GHz–3.4GHz	89W
Intel Pentium 4 2.2GHz–2.4GHz	80W–90W
Intel Pentium 4 2.4GHz–3.0GHz	90W–105W
Intel Celeron Socket 478	45W–65W
ATX motherboard	40W–65W
PC133 RAM	12W
PC2100+ DDR RAM	10W
PC3200+ DDR2 RAM	7.5W
CD-ROM drive	20W
CD-RW drive	30W
DVD-ROM drive	25W
5,400RPM IDE hard drive	15W
7,200RPM IDE hard drive	25W
Floppy drive	5W
Network card or modem	5W
Sound card	7–18W
SCSI controller card	25W
FireWire or USB 2.0 controller card	40W
USB device	5W
FireWire device	8W
CPU or case fan	2W

If you are replacing a power supply in a system that experienced a failure, then you can get a replacement that provides the same amount of power or more; but, if expect that your system has more devices than the power supply can support or if you are building a custom system, you can follow these steps:

1. Locate on Table 6-1 any devices that your system is going to contain and record the power usage.
2. Add up all of the wattages and divide your total by 80% (0.8) to give your system some extra room to grow and to accommodate any items you may have missed.
3. If you did not get a number that matches an available power supply rating, then choose a power supply with the next highest power rating.

Power Connectors

A computer uses many different types of power connectors. They include the main power, floppy, peripheral, 12 volt (+12V2DC), and SATA connectors.

The ATX main power connector

On new power supplies, the ATX main power connector is a 24-pin Molex 39-01-2240 connector, and, in many cases, it is a 20+4-pin connector for backward compatibility with older version 1.3 motherboards. This is a keyed connector, so you can't connect it the wrong way without a lot of effort. You can see this and other connectors in Figure 6-1.




Figure 6-1:
Power connectors that may be used on power supplies.

Table 6-2 provides you with the pin configuration and voltage supply for the 24-pin Molex main power connector. When troubleshooting faulty power supplies, you will need to know what power is supposed to be supplied by each pin in the power connector so that you can identify that the issue you are troubleshooting is related to a problem of power supply.



Remember that the main power supply connector provides power at 3.3 VDC, 5 VDC, and 12VDC.

Table 6-2 Pin Configuration for the Main Power Connector

<i>Pin</i>	<i>Signal</i>	<i>Color</i>
1	+3.3VDC	Orange
2	+3.3VDC	Orange
3	COM	Black
4	+5VDC	Red
5	COM	Black
6	+5VDC	Red
7	COM	Black
8	PWR_OK	Gray
9	+5VSB	Purple
10	+12V1DC	Yellow
11	+12V1DC	Yellow
12	+3.3VDC	Orange
13	+3.3VDC	Orange
14	-12VDC	Blue
15	COM	Black
16	PS_ON#	Green
17	COM	Black
18	COM	Black
19	COM	Black
20	Reserved	Varies
21	+5VDC	Red
22	+5VDC	Red
23	+5VDC	Red
24	COM	Black

The floppy drive connector

Although floppy drives are becoming nearly obsolete, the floppy connector is an AMP 171822-4 four-pin connector and has its pin configuration listed in

Table 6-3. This connector provides power to a floppy drive at 5V, but is capable of supplying power at either 5V or 12V to support devices from manufacturers who decide to use this connector.

Table 6-3 Pin Configuration for the Floppy Power Connector

<i>Pin</i>	<i>Signal</i>	<i>Color</i>
1	+5VDC	Red
2	COM	Black
3	COM	Black
4	+12V1DC	Yellow

The peripheral connector

Until SATA drives rule the world, the four-pin peripheral connector is likely the most used power connector on most computers because it is used for hard drives and optical drives. This connector is an AMP 1-480424-0 or Molex 8981-04P, and its pin configuration can be found in Table 6-4. Like the floppy connector, it provides power at both 5V and 12V, but the hard drives draw their power at 12V, rather than the 5V which is sufficient for floppy drives. Hard drives require more power to get their heavier disk spinning.

Table 6-4 Pin Configuration for the Peripheral Power Connector

<i>Pin</i>	<i>Signal</i>	<i>Color</i>
1	+12V1DC	Yellow
2	COM	Black
3	COM	Black
4	+5VDC	Red

The 12-volt power connector

The 12-volt power connector (+12V2DC) is also called a P4 connector and is used to supply additional power directly to the processor. It uses a Molex 39-01-2040 connector and its pin configuration is listed in Table 6-5.

Table 6-5 Pin Configuration for the +12V2DC Power Connector

<i>Pin</i>	<i>Signal</i>	<i>Color</i>
1	COM	Black
2	COM	Black
3	+12V2DC	Yellow
4	+12V2DC	Yellow

The Serial ATA connector

The new Serial ATA (SATA) drives add two new connectors to the long list that are used by computers. One of the connectors is used for data, and the other is used for power. The data connector is discussed in Book III, Chapter 1. The power connector uses a Molex 88751 connector, and an additional wire (5) supplies 3.3V to devices that require it (most functions of the SATA drive run on 12 volts). The SATA connector has its five wires connected to 15 pins on the connector. The SATA connector pin configuration is listed in Table 6-6.

Table 6-6 Pin Configuration for the Serial ATA Power Connector

Wire	Signal	Color
1	+12 V1DC	Yellow
2	COM	Black
3	+5 VDC	Red
4	COM	Black
5	+3.3 VDC	Orange

Power Supply Form Factors

Some other form factors also provide power for ATX motherboards. They exist mainly to fit in different sizes and styles of cases. These power supplies typically have the same types of power connectors that were detailed in the previous sections. The specifications for the power supplies usually define mounting points and the maximum footprint or space that the power supply may occupy. In addition, the specifications also state whether cooling fans are required. In addition to the most common form factor, which is ATX v2.2, the form factors for the power supplies include the following:

- ◆ **CFX v1.3** is the *Compact Form Factor*, typically supplying power in the range of 220W to 300W.
- ◆ **LFX v1.1** is the *Low Profile Form Factor*, typically supplying power in the range of 180W to 260W.
- ◆ **SFX v3.1** is the *Small Form Factor*, typically supplying power in the range of 160W to 300W.
- ◆ **TFX v2.2** is the *Thin Form Factor*, typically supplying power in the range of 180W to 300W.

The size specifications for these, as well as the ATX power supply, are shown in Figure 6-2. This figure also shows the mounting holes for each type of power supply. Computer power supplies are held in by three or four screws

and are usually positioned in the top-rear portion of the computer. After removing the power connector from the internal devices, you can remove the screws and remove the power supply.



Take care when removing the screws. Some cases have rails that support the power supply, but with others, the power supply could drop onto the computer components in the case when the last screw is removed.




Figure 6-2:
Power supply unit form factors.

Using AC Adapters

While desktops have large internal power supplies to get AC building power converted to 3.3V, 5V, and 12V for internal use, laptops typically do not have the internal space to support such a large power supply. Most laptop AC adapters are external and have the job of converting the building power to a voltage that is required by the laptop. Laptops typically require between 12V and 20VDC. The laptop then has internal components that divide this incoming voltage into segments that are appropriated for the internal devices, such as 12V, 5V, and 3.3V. This allows for part of the power management components to exist outside of the laptop.

What about DC adapters?

DC adapters allow you to convert DC voltages to different DC voltages, changing the power characteristics, such as current, to allow 12V power from a vehicle to be used for computers and other equipment that does not use a 12V

power directly. You may use these from time to time to power other devices, but, as the power requirements for most laptops are high, you may not be able to use a 12VDC adapter to provide it power.

In some cases, the power supply may use generic output voltages and connectors, while, in other cases, the voltage or connectors may be customized by a manufacturer, meaning that replacement power supplies are tied specifically to that manufacturer. The AC adapter in Figure 6-3 has a fairly standard connector, and, if you look closely, you can see that the output voltage is 18.5 volts.



Figure 6-3:
An AC
power
adapter for
an HP
laptop.

Output voltage



AC adapters are used with a wide variety of devices that you deal with every day, converting AC power to DC power for cell phones, cordless phones, PDAs, and a host of other devices. You should always take care to match the correct power adapter to the power requirements for the devices to ensure that you don't damage the device — or worse.

When working with your AC adapter, you should always keep your eye on damage that may have occurred along the length of the cord from wear or perhaps pets. Neat people actually tend to be harder on adapter cables than disorganized ones because neatniks tie the cables up every time they pack their laptops, consolidating the weight of the cable mass right next to the brick of the adapter, causing the cable to wear quickly in the section nearest the adapter. To avoid this problem, the cable should be coiled *loosely* and placed in your carrying case.

Working with UPS and Suppressors

Some people have such good and stable power that they take it for granted. But this is not the case for many of us, and we live with a variety of power issues such as these:

- ◆ **Spikes:** Sharp increases in line voltage caused by a problem with the power utility equipment, which may be traced back to faulty transformers or lightning strikes.
- ◆ **Power surges:** Similar to spikes, power surges are increases in line voltage that usually last longer than a spike.
- ◆ **Noise:** Interference caused by items not directly connected to the power system. Power lines running near florescent lights or microwaves can have their AC sine waves altered or chopped, which affects devices receiving their power.
- ◆ **Blackout:** A total interruption in the line power. The ATX power supply specification has a minimum voltage hold-up time of 17ms, or about 1/60 of a second. This is the amount of time for which the power supply can sustain computer operation during a power interruption; in other words, if your lights flick on and off due to a very short power outage, your computer may continue to run over that outage, but it would have to be a very brief power outage.
- ◆ **Brownout:** The opposite of a power surge in that the line power is below normal for a period of time ranging from seconds to minutes. If the brownout exceeds the Input Under Voltage limits of the power supply, the power supply shuts off.

Three basic devices help you deal with these power situations:

- ◆ **Surge suppressors** protect your equipment against spikes and surges by quickly terminating the power connection before your equipment is damaged.
- ◆ **Line conditioners** take that a step further by cleaning the noise on the line and dealing with some under- and over-voltage situations, as well as generating alarms to notify you of any situations.

In most cases, line conditioners deal with surges and spikes while letting your equipment continue to function.


- ◆ **Uninterruptible Power Supplies (UPS)** take power protection beyond line conditioners by keeping your equipment working even after total power failures.

There are two main types of UPS:

- ◆ **Standby**, also called *offline*, does not protect you from surges or spikes but takes over quickly in the event of voltage drops or fails. UPS units that are marketed toward home users are usually standby units.
- ◆ **Inline**, also called *online*, UPS always operate between the line and your equipment. In addition to providing power, when the UPS battery power is below a specific level, they can initiate a graceful shutdown of the devices that are connected to it.


Most UPS include scales or lights to indicate the current load on the UPS as a percentage of its maximum capacity, as well the battery charge condition. It is important to either check these indicators regularly or install software that monitors the UPS for you. UPS indicators indicate capacity levels as well as error conditions.

UPS often have a variety options for plug types to supply devices with power. A standard server UPS is shown in Figure 6-4.



UPS front view

Figure 6-4:
Front and
rear of a
standard
server UPS.



UPS rear view

Most UPS are rated in volt-amps (VA), which indicate how much power they can supply. VA ranges for UPS start around 200VA and go up to tens of thousands of volt-amps for corporate data centers.



In a perfect world, watts and volt-amps would be identical. Early power supplies were *capacitor input* power supplies that used capacitors to even out the power supplied by a building's AC power. The method of converting power creates a variance between watts and VA of 55% to 75%, which is called the Power Factor (PF). This gives us a formula of $watts = VA * PF$. Capacitor input power supplies are still in use on most inexpensive power

supplies. The solution is the *Power Factor Corrected (PFC)* power supply, which uses watt and VA ratings that are the same. Most high-end equipment that was manufactured since 1996 should include a PFC power supply.

Before buying a UPS, total the power requirements for all of the devices you want to plug into it, as well as the amount of time for which you need protection. The protection time is usually between 10 and 20 minutes; longer periods are possible if you want to pay the additional cost. Use the information to make sure you get a UPS that is powerful enough to protect all your expensive gear.

Getting an A+

This chapter sheds light on power and how it interacts with your computer. The following points are covered:

- ◆ The current specification of power supplies is ATX, and the ATX power supply will be found in one form factor or another on most current computers.
- ◆ The main power connector has 20 or 24 pins, four-pin floppy and peripheral connectors supply 5 volts and 12 volts, and five-pin SATA connectors supply 3.3, 5, and 12 volts.
- ◆ Power supplies convert 120VAC or 240VAC power into DC power for the computer.
- ◆ Power supplies come in a variety of power sizes, from 160 watts up to and beyond 800 watts.
- ◆ AC adapters are used on laptops to convert power from AC to DC and usually automatically switch between 120V and 240V.
- ◆ Surge suppressors, line conditioners, and UPS deal with various power-supply problems, including surges and brownouts.

Prep Test

- 1** What are the main power voltages that are provided by standard computer power supplies? (Select all that apply.)
 - A 2.1V
 - B 3.5V
 - C 5V
 - D 12V

- 2** What is the name given to the main type of power supply used in modern computers?
 - A AT
 - B NLX
 - C ATX
 - D WPA

- 3** What components influence the purchase decision of a power supply? (Select all that apply.)
 - A Hard drive
 - B Monitor
 - C RAM
 - D Parallel port scanner

- 4** What are the standard input voltages for power supplies? (Select all that apply.)
 - A 110VDC
 - B 120VAC
 - C 240VDC
 - D 240VAC

- 5** What is the main function of an AC adapter for a laptop or other device?
 - A Reduce heat involved in converting power between systems.
 - B Supply 3.3V, 5V, and 12VDC power to your computer through a 20- or 24-pin connector.
 - C Provide power availability during brownout and blackout conditions.
 - D Convert AC power into DC power.

Working with Power

- 6** Which of the following power conditions does a line conditioner not offer protection against?
- A Blackout
 - B Brownout
 - C Spike
 - D Surge
- 7** Which of the following power conditions does a UPS not offer protection against?
- A Surge
 - B Blackout
 - C Line noise
 - D Energy Star

Answers

- 1** **B, C, D.** ATX-based power supplies provide power to devices at 3.3V, 5V, and 12VDC. *See "Identifying the Purpose of Power Supply Units."*
- 2** **C.** Most current computers get power from an ATX power supply. *Review "Identifying Power Supplies."*
- 3** **A, C.** Both of these devices, as well as most other internal devices, affect the total power required by a power supply. Monitors have their own internal power supplies, and most parallel port scanners come with their own AC adapters. *Check out "Identifying Power Supplies."*
- 4** **B, D.** 120VAC and 240VAC are the two power input voltages that are handled by power supplies. *Peruse "Identifying the Purpose of Power Supply Units."*
- 5** **D.** The main function of an AC adapter is to convert AC power into DC power for use in a laptop. Power is provided to the laptop through a connector of some type, but it will be a single voltage and will not be a 20- or 24-pin connector. A UPS would be required to provide power during a blackout. *Take a look at "Using AC adapters."*
- 6** **A.** Line conditioners are able to deal with some reduced voltage situations, depending on their severity, but are not able to deal with complete blackouts. *Peek at "Working with UPS and Suppressors."*
- 7** **D.** Rapid cutover is a manner of switching to UPS power when using a standby UPS, while the other items are all power line problems that a UPS provides solutions for. *Look over "Working with UPS and Suppressors."*

Book III

Outside the Box

The 5th Wave

By Rich Tennant



"Here's a little tip on disassembly that you won't find on the A+Certification test."

Contents at a Glance

Chapter 1: Ports, Cables, and Connectors	291
Chapter 2: Installing and Configuring Input Devices	315
Chapter 3: Installing and Configuring Output Devices	335
Chapter 4: Examining System Resources	351
Chapter 5: Managing Printers	379
Chapter 6: Working with Multimedia Devices	407
Chapter 7: Dealing with Portable Computers	421

Chapter 1: Ports, Cables, and Connectors

Exam Objectives

- ✓ Understanding fundamental principles of using a personal computer
- ✓ Identifying ports and cables

In this chapter, you examine the ports on the back of the computer as well as the cables and connectors that are used to attach devices. In other chapters of this book, you examine different devices, and this chapter looks at getting them all connected. The other devices and chapters include:

- ◆ Printers in Book III, Chapter 5
- ◆ Internal drives in Book II, Chapter 5
- ◆ Monitors in Book III, Chapter 3
- ◆ USB devices in Book II, Chapter 1
- ◆ FireWire devices in Book II, Chapter 1
- ◆ Modems in Book III, Chapter 2

A CompTIA A+ Certified Professional must be able to identify various types of cables and connectors. This chapter introduces you to the most common cables and connectors that you will encounter when using computer systems.

Identifying Common Computer Ports

One of the first things to look at are the different types of ports on the back of your computer. *Ports* are connection points on your computer that allow for devices to be connected to your computer. The connection point will have a connector that accepts a cable with a matching connector. Knowing what the ports are used for is important when connecting devices to your computer. You may already be familiar with some of the common ports, such as serial, parallel, and USB. In this section, I discuss these ports as well as IEEE-1394 ports, which are often called FireWire.

Serial and parallel ports

Two of the most common computer ports are the serial ports and the parallel ports. Most computers have one parallel port and two serial ports. These ports are used to connect devices such as modems and printers to your computer.

Parallel ports send data over multiple wires simultaneously, while *serial ports* send data over only one wire at a time. Because parallel communication allows for multiple streams of data, it provides higher data transfer rates than serial communication.

Parallel ports were implemented on the personal computer when it was introduced by IBM in 1981. They used nine-wire cables to connect two devices. This enables the cables, at any given time, to deliver eight bits of data, with the extra wire being used as a ground. However, because there is no way to accurately control the flow of the signal down each one of the wires in the cable, it is recommended that the length of the cable be less than six feet (about 2 meters).

Controlling the flow of data became more of a factor when *bidirectional communication* (sending and receiving data with the device) was implemented over the parallel port. The standard for bidirectional communication was delivered in 1994 in the IEEE-1284 specification, which allows for high-speed, two-way communication over the parallel port. This also opened two new specifications for the port: the *Enhanced Parallel Port (EPP)* and the *Extended Capabilities Port (ECP)*. The maximum cable length for an IEEE-1284 cable is about 30 feet (10 meters).

The EPP-type parallel port is used primarily for non-printer peripherals, while the ECP-type parallel port is designed to accommodate new high-speed printers and scanners. In order to better handle high-speed data communication, the ECP-type parallel port also implements the use of a *DMA (Dynamic Memory Access)* channel; this grants the port direct access to a section of RAM.

Serial ports, on the other hand, deliver data sequentially down a single wire. Eight bits of digital data are converted into an analog signal by using a system called *baud*. *Baud rate* refers to the number of state changes (tones) that are made on the wire in any given second. If you have heard a Fax machine or modem making a connection, you have heard the state changes as a series of squeals. Baud rate is very different from bits per second (bps), which measures the amount of data that is transferred. At one point, 300 bps modems communicated at 300 baud, but compression standards adopted by the communications industry allowed more data to be delivered at the same baud rate. Today, 56Kbps (57,344 bps) modems communicate at 9,600 baud.

Depending on the baud rate used to transfer data, the length of the cable can range up to 3,000 feet (just over 900 meters). For data transfer rates at 9,600 baud, the maximum cabling length is 250 feet (just over 75 meters). The *RS-232C standard*, which is used as a basis of serial communication, recommends a maximum cable length of 50 feet (about 15 meters).



Find out how to troubleshoot serial and parallel port problems by reading Book IV Chapter 2.

Universal Serial Bus (USB)

USB is a new method of communicating using serial communication methodologies. The standards for USB 1.0 were released in 1996, USB 1.1 in 1998, and USB 2.0 in 2000. The goal of USB was to revolutionize the way serial communication was conducted. In an effort to fulfill this goal, USB uses a new cabling system that allows up to 127 devices to be connected together. It also delivers power to the devices that are connected on this bus. To go along with this new cabling, the specification for USB dictates that all devices should support Plug and Play. With these two pieces of the puzzle put together, USB enables you to plug devices in and have them work without having to worry about power cables or drivers. Some USB devices that require a large amount of power may use a supplemental power supply, but the USB bus will power most devices.


USB 2.0 increased the fastest transfer rate for USB 1.0 from 12 Mbps to 480 Mbps. This has allowed it to be used for devices that require faster transfer rates, such as hard drives and CD-ROMs. For more information about USB, see Book III, Chapter 3.



USB has had a performance boost up to 480 Mbps with the introduction of version 2 of the USB standard. For the exam, you need to pay attention to the version of USB that is being covered in the question. If it is USB 1.0, then you will still be limited to the 12 Mbps transfer rate.

FireWire (IEEE-1394)

FireWire is an Apple trademark for the IEEE-1394 standard. The 1394 standard implements a version of serial communication across a wiring network that is similar to USB. IEEE-1394 enables the connection of 64 devices on a bus that supports 50 to 400 Mbps, and with the release of FireWire 800, the transfer rates have approached 800 Mbps. One of the goals of the IEEE-1394 standard was to replace SCSI, which is covered in Book II, Chapter 5.



Don't confuse IEEE-1394 (FireWire) with IEEE-1284 (mentioned in the "Serial and parallel ports" section, earlier in the chapter), which deals with bidirectional communication over the parallel port.

Although some disk drives have been implemented through the IEEE-1394 standard, it has seen the most growth in the area of data, voice, and video. While many PCs have been shipping for the last few years with USB ports, IEEE-1394 has not seen the same level of adoption by the PC manufacturing industry. However, many device manufacturers are supporting IEEE-1394 in their devices, and interface cards are readily available to be added to your PC.

FireWire is covered in Book II, Chapter 1.

Keyboard

Because you will certainly want to enter data into your computer, the keyboard connector on the back of your computer plays an important role. There are two traditional types of keyboard connectors: P/S2 (or Mini-DIN 6) and DIN 5. The mini-DIN connector is the standard keyboard connector on computers; but it faces being supplanted by USB.

For details on connecting keyboards to your computer, review Book III, Chapter 2.

Monitor

You also need to see what you're doing; so the monitor connector also plays an important role. With the demise of the MDA (Monochrome Display Adapter), HGC (Hercules Graphics Card), CGA (Color Graphics Adapter) and EGA (Enhanced Graphics Adapter) monitor types and the release of VGA (Video Graphics Array) and SVGA (Super Video Graphics Array) monitors, the monitor connector on the back of your computer changed from a DB-9 female connector to an HD (high density) DB-15 female connector. With the high adoption rate of LCD monitors, a new port is showing up on computers and video cards, supporting the higher throughput and digital signals that LCD monitors require. This new port is a DVI (Digital Visual Interface) port and accepts a monitor connector that has up to 24 digital pins and optionally five analog pins (for backward compatibility).

You can read more about connecting monitors to your computer in Book III, Chapter 3.

Comparing Cable Types

In this section, you examine several different types of cables and discover the uses of each type. Some of the cables that you will see are used inside of your computer, while others are external cables. Once you have completed this section, you should be able to identify the basic types of cables and where they are used.

Ribbon

Ribbon cables are often used to connect components inside a computer, such as hard disk drives and floppy disk drives. They are made up of several wires that are laid out parallel to each other in such a way that they resemble a ribbon, as shown in Figure 1-1. These cables usually have keyed connectors on them that prevent them from being incorrectly connected to devices. You should note the small tab halfway down the edge of the black connector. This tab matches a groove on the device that it is attached to.




Figure 1-1: Ribbon cables are usually found on the inside of your computer.

Small tabs

You will read more about ribbon cables when connecting storage devices inside of your computer, which is covered in Book II, Chapter 5.

Twisted pair

Twisted pair cables consist of three or four pairs of wires. The grading level is based on how the wires are arranged inside of the cable, rather than the number of wires. In Figure 1-2, you can see how each pair of wires is twisted together at a specific rate, and then how all of the pairs are twisted together. This procedure reduces the effect of cross talk or interference between the pairs of wires and interference from external sources. The differences between the different grades of cable include the quality of production material and the overall number of twists per pair of wires per foot of cable. Most networking is done with unshielded twisted pair (UTP) cable, which is more

susceptible to interference; but is more flexible and therefore easier to work with than shielded twisted pair (STP) cable, which has a metal shielding over each pair of wires.

Figure 1-2:
How the wires are twisted is one of the differences between the categories of cable.

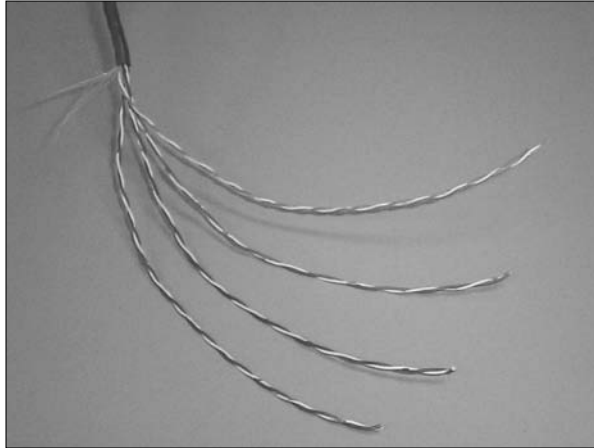


Table 1-1 provides a brief listing of the different types of twisted pair cables and their uses.

Table 1-1		Types of Twisted Pair Cabling	
<i>Type</i>		<i>Usage</i>	
Category 1 (CAT1)		Unshielded twisted pair designed for transmission of audio signals. This wiring is intended to be used for transferring signals to speakers and other devices.	
Category 2 (CAT2)		Unshielded twisted pair used for low-speed transmissions. This wiring was designed for analog telephone applications. CAT2 cables are also capable of carrying lower-speed networking data, as in the case of 4 Mbps token ring networking.	
Category 3 (CAT3)		Unshielded twisted pair cable that is designed with data networking in mind. This category is similar in design to all of the newer categories, but some electrical characteristics are unique to each. CAT3 cables carry 10 Mbps Ethernet data.	
Category 4 (CAT4)		Unshielded twisted pair cable that was designed to carry 16 Mbps token ring data.	

<i>Type</i>	<i>Usage</i>
Category 5 (CAT5)	Unshielded twisted pair cable that was designed to carry 100 Mbps Ethernet data. It also meets the requirements for 1000Base-T. <i>CAT5e</i> is an enhanced version of the cable that better meets the standards of 1000Base-T by reducing crosstalk.
Category 6 (CAT6)	A proposed standard for unshielded twisted pair cable that supports frequencies of 250 MHz Ethernet data. <i>CAT6a</i> increases the frequencies that the cable supports and allows for data transfers rates up to 10 Gbps.
Category 7 (CAT7)	A proposed standard for unshielded twisted pair cable that supports frequencies of up to 600 MHz.

You can read more about networking cables in Book VIII, Chapter 1.



Pay close attention to the uses of CAT3, CAT4, and CAT5 cable types because these map directly to network specifications. You will be tested on the minimum required cable type for 10 Mbps Ethernet and 100 Mbps Ethernet.

Thick and thin coax

Thick coax is rigid half-inch coaxial cable that is capable of carrying 10 Mbps Ethernet data a distance of 500 meters. Thick coax is also referred to as *Thicknet* or *10Base5*.

Thin coax is more flexible quarter-inch coaxial cable that is capable of carrying 10 Mbps Ethernet data a distance of 185 meters. Thin coax is also referred to as *Thinnet* or *10Base2*.

For more information about thick and thin coax use in computer networking, flip to Book VIII, Chapter 1.

Fiber

Fiber optic cable is so thin that it is measured in microns (>m), or millionths of a meter (or, put another way, thousandths of a millimeter). Fiber optic cable for data networks comes in two core diameter thicknesses 50>m and 62.5>m. There are also cables that support a single communication mode or multiple communication modes, which refers to how light traverses the cable. Multimode cables can support distances of up to 1 km, while single-mode cable can support distances of up to 140 km.

Fiber optic cable is also covered in Book VIII, Chapter 3.

Cable Orientation

Cable orientation refers to the way two connectors fit together, so it could actually be referred to a connector orientation. More to the point, it is how a connector on a cable attaches to another connector, which is usually on a computer or motherboard. Just about every time the term cable orientation is used, it will be talking about the internal ribbon cables used to connect serial ports, parallel ports, floppy drives, hard drives, and CD-ROM drives to a motherboard. Cable orientation is not usually used when discussing external connectors, which are used on the outside of your computer and tend to be able to connect only one way. The pins or holes on a connector are numbered, and the hole numbered one on a cable connector needs to match pin number one on the computer's connector.



The connector with the pins is called the male connector, and the connector with the holes is called a female connector. Even though the female connector does not have pins, it is typical to refer to the holes as pins on the female connector.

Internal cables are used on the inside of the computer. Some devices use unique cables, such as your sound card and CD-ROM drive. In many cases, these unique cables have molex connectors (molded plastic connectors) that are usually keyed to prevent mistakes on connecting. A keyed connector is designed to connect only one way. You also see molex connectors on the power leads that connect your power supply to devices such as hard drives and the motherboard.



Remembering one simple rule can prevent most connection errors. On a ribbon cable, Wire 1 connected to Pin 1 has the colored (usually red) stripe on it. This makes it easy to spot. On most devices that you are connecting the cable to, Pin 1 will be the pin closest to the power connector.

Figure 1-3 shows two IDE connectors on the motherboard. Note that these are keyed and also have a small arrow pointing to Pin 1. Most of the male connectors on the motherboard have a wall around the pins that has a slot, or key opening that prevent errors, as long as the cable you are using is also keyed. Keyed connectors on your cards or motherboard have a small notch cut from the side of the connector, which is in the middle of the long side of the connector in Figure 1-3, whereas the keyed cables have a small lump or tab on them, which can be seen on the long edge of the connector (refer to Figure 1-1). The only way for the cable to be plugged into the connector is if the tab matches the notch, and there is only one way for that to happen, which will line up Pin 1 on your connectors.




Figure 1-3:
Pin 1 is usually identified by some marking.

Arrows indicating Pin1

Connector Types

Many, many types of cables exist on the market these days: video cables, modem cables, printer cables, and extension cables, to name a few. Even though there are many types of cables, there are very few types of connectors. The following sections look at the most common types of connectors that you will encounter.

IBM Type 1 Connector

The IBM Type I connector is commonly used with IBM Token Ring networks, and will either connect the workstation to a hub or connect multiple hubs together. Now, unlike typical connectors that are either male (with pins) or female (with holes) and must be paired with a connector of the opposite gender to allow them to be connected, the Type I connector can be connected to any other connector. IBM has called the connector a hermaphrodite, being composed of both genders of connectors. Figure 1-4 shows a Type I connector, with two connectors joined at the top of the figure.




Figure 1-4:
The IBM
Type 1 data
connector is
a herma-
phrodite.

DB-9

The DB-9 connector is a D-shell-type connector. They're called *D-shell* connectors because of their shape. The DB-9 has nine connection points arranged in two rows. This connector used to be found on the back of your computer in its female form, to be used with a CGA or EGA monitor. Now, you'll find a male DB-9 connector on the back of your computer as one of the two types of serial connections for your COM ports. Figure 1-5 shows the male and female DB-9 connectors.

DB-15

DB-15, also known as HD DB-15, has the familiar shape of the D-shell connector but contains a higher density (HD) pin arrangement, with three rows of five connection points rather than the traditional two rows. It is usually used for VGA and SVGA monitor connections and should be found on your computer in its female form. Figure 1-6 shows the DB-15 connector.

DB-25

Once again, you see the familiar shape of the D-shell connector on a DB-25, but this time you get 25 connections arranged in two rows. This connector is found on your computer in its male form in the role of a serial port and in its

female form as a parallel port. You will also find the 25-pin male connector on the back of most modems. Figure 1-7 shows a parallel cable and a serial 9-25 pin converter, which is used to convert the 25-pin serial connector on the back of a computer, down to a 9-pin port to be compatible with a 9-pin serial cable.





Figure 1-5:
The DB-9
connector.



Figure 1-6:
The high
density HD
DB-15.

Figure 1-7:
The DB-25
connector.




Serial 9-25 pin converter

Parallel cable

Centronics 36 and 50

A parallel cable has a DB-25 connector at one end and a 36-pin male Centronics connector at the other. You will find the female Centronics receptacle on your printer. Also, you will often see 50-pin Centronics connectors used with SCSI equipment, such as external disk or tape drives. Figure 1-8 shows 36-pin and 50-pin Centronics connectors.

Figure 1-8:
Centronics
36 and 50
connectors.



50-pin

36-pin

RJ-11

RJ is short for *registered jack*, and it has small modular connectors that clip into matching holes. The RJ-11 connector is a standard modular connector that is used for telephones. It accepts four wires, usually in the form of a flat cable, rather than twisted pair cable. Analog telephone service is usually only on the two middle wires. You should see a female portion of this connector on your modem.

RJ-45

The RJ-45 connector is the larger cousin of the RJ-11 connector. They usually accept eight wires — or four pairs. This connector is used for 10BaseT, 10BaseTX, and token ring networking. (Basically, the RJ-45 is used anywhere UTP cables are used.) Officially, the RJ-45 connector was designed for voice-grade circuits, and what is now referred to as the RJ-45 is officially known as an *8-pin connector*.



RJ-11 connectors should not be confused with RJ-45 connectors. Figure 1-9 shows a RJ-11 connector on the left and a RJ-45 on the right.




Figure 1-9:
RJ-11
connectors
are smaller
than RJ-45
connectors.

BNC

The *BNC* (*Bayonet Neill-Concelman*) is also sometimes referred to as *British Naval Connector* and a number of other names, but the actual name is a moot point for the exam. This connector is used for Thinnet networking (see “Thick and thin coax,” earlier in this chapter). BNC connectors join together with a male BNC that has a protruding point that fits into the female connector, and then a ring is turned on the outside edge to lock the connectors together in

much the same way that a bayonet attaches to a rifle (hence, the name). Figure 1-10 shows a BNC T-connector and a network card with a BNC connector.




Figure 1-10:
Thinnet
Ethernet
uses BNC
connectors.

BNC T-connectors are used to attach two Thinnet cables to a network card on the back of your computer. The T-connector has a male connector to attach to your network card and then two female connectors to accept the two Thinnet cables. A Thinnet network segment makes one long chain out of all of the computers that are on the segment by using T-connectors.

PS/2 or Mini-DIN 6

The PS/2 (also known as the Mini-DIN 6) connector is usually used for keyboards and mice for AT and ATX computers. This connector is now favored over the larger DIN 5 connector. Figure 1-11 provides a sample of the Mini-DIN connector.

Universal Serial Bus (USB) connectors

USB 1.0 and USB 2.0 both use proprietary connectors (shown in Figure 1-12) to connect up to 127 devices together. All of these devices must be hot swappable, or able to be added or removed without turning the computer off, thanks to the specification for the standard. There are two main types of connectors in USB: Type A for hosts and Type B for devices.

Figure 1-11:
PS/2
connectors
are used for
keyboards
and mice.



Figure 1-12:
The USB
connector is
unique to
the USB
system.




Type A

Type B

IEEE-1394 (FireWire) connectors

FireWire is an Apple Computer trademark for devices that match the IEEE-1394 specification. This specification allows 63 devices (without hubs) to be connected together in a single bus. It also uses a proprietary connector that is similar in size to the USB connector but has a different shape, as seen in Figure 1-13.

Figure 1-13:
The
FireWire
connector is
unique to
the FireWire
system.



Standard External Cables

External cables are used on the outside of the computer to connect peripheral devices. The most commonly used connector on external cables is the D-shell connector (named for its shape), which is usually called *DB*. The numbering of pins on male connectors starts with the top-left connector if you are looking at the connector with the D pointing toward the ground. With the connector in this orientation, the upper-left pin is Pin 1, and the order goes across the connector and left to right for all subsequent rows. If you are dealing with a female receptor, then Hole 1 or Pin 1 is in the top-right of the connector if the D is pointing down, and the other pins are numbered right to left. This allows Pin 1 on a male connector to match Pin 1 on a female connector. This section will look at the three most common cables.

Parallel Cable

The parallel cable will have a male DB-25 connector on the end that connects to your computer, and a male Centronics 36 connector on the end that connects to the printer. This style of cable will either be a standard parallel cable or an IEEE-1284 Bi-directional cable. If the cable is the latter, you should expect to see IEEE-1284 labeled on the cable or connector.

In addition to this cable, you may also see a parallel extension cable, which will have a male DB-25 connector to attach to your computer and a female DB-25 connector to attach to a parallel printer cable. This cable could easily be confused with a serial cable for a computer with a 25-pin serial connector. The only way to positively identify this cable is to use a multi-meter or a continuity tester to see what pins on one end of the cable match which holes on the other end. For a parallel extension cable, this will be a straight-through cable, and each pin, 1 through 25, will match with the appropriate hole on the other end of the cable.

Serial Cable

Numbering the pins on the connectors makes it easier to describe what the cable looks like at the connector level. This may make more sense when you look at the pin configuration for a standard 9-pin to 25-pin serial modem cable, and the 9-pin to 9-pin null modem cable in the next section. Modem cables are designed for a computer with either a male 9-pin connector or a male 25-pin connector, so the cable has either a female 9-pin or female 25-pin connector on the end that connects to the computer, and a 25-pin male connector on the end that connects to the modem. Most modern computers will have a 9-pin serial connector, if they have a serial port at all. The pin configuration and function for a 9-pin to 25-pin serial cable are summarized in Table 1-2.

Table 1-2 Serial Cable Pin Connections

<i>Computer (9 pin)</i>	<i>Data Direction</i>	<i>Modem (25 pin)</i>	<i>Role</i>
3	→	2	Transmitted Data (TD)
2	←	3	Received Data (RD)
7	→	4	Ready To Send (RTS)
8	←	5	Clear To Send (CTS)
6	→	6	Data Set Ready (DSR)
1	→	8	Data Carrier Detected or tone from a modem (DCD)
4	→	20	Data Terminal Ready (DTR)
5	–	7	Signal Ground (SG)
9	←	22	Ring Indicator or ringing tone detected (RI)

Null Modem Cable

You may need, at times, to transfer data between two computers that are close to each other by using their serial ports. To do this, you will use the same communications programs that you would use if you were actually using modems and telephone lines to connect the computers. This serial communication that bypasses a modem makes use of a null modem cable. The null modem cable nullifies or eliminates the need for a modem by directly connecting the send and receive ports of the devices on either side of the cable. Table 1-3 shows how the pin connections are made between devices. If you look at the roles of the pins on each side of the connection, you see that each pin that sends data is connected to the receive pin for that data on the other side of the cable, with the result that a transmit data pin is connected to a receive data pin.

<i>Role</i>	<i>Computer (9 pin)</i>	<i>Data Direction</i>	<i>Computer (9 pin)</i>	<i>Role</i>
RD	2	->	3	TD
TD	3	<-	2	RD
RTS	7	->	8	CTS
CTS	8	<-	7	RTS
DCD and DSR	1 and 6	<-	4	DTR
DTR	4	->	1 and 6	DCD and DSR
SG	5	--	5	SG

Null modem cables are also used to configure many network devices, such as switches and routers, using a console port on the network device and a terminal communication program.

Viewing Cable Adapters

Many types of adapters can join different types of cables and devices together. The following sections give a brief description of some of the major types.


Barrel connectors

The term *barrel connectors* comes from Thinnet networking, where cables could be extended by means of an adapter, resembling a small barrel, that accepted a male connector on either end. This term is now often applied to any connector that extends the length of a cable by joining two cables together but that does not change the pin configuration. Figure 1-14 shows a barrel or cable connector for an RJ-45 cable.

Gender changers


Gender changers are a bit of a strange beast. They are straight-through connectors that don't change the order or connection of the pins; they only change the sex of the connector that they are attached to. They look like a connector that has either two male or two female ends (as shown Figure 1-15), and when attached to a cable of one sex, the connector becomes the other. These are used only in rare cases where your cables have the wrong sex connector, which is a situation that I often encounter with video connectors when working with KVM (Keyboard, Video, and Mouse) switches and cables that are made by different manufactures.

Figure 1-14: Barrel connectors let you extend the length of cables.



Barrel connector

Figure 1-15: Gender changers change the sex of a connector.



Gender changer

Null modem

Null modem cables act as a replacement for two modems communicating with each other. If you have two computers that you want to connect directly together, then you can attach two modems and have one dial the other and transfer the files, or you can use a null modem cable. The *null modem* removes

the modems from the equation and connects the send pins on one serial port to the receive pins on the other, and vice versa. If you don't have a null modem cable, you can use a regular modem cable and a null modem adapter. The null modem adapter looks like a gender changer, except that it changes the pin configuration as well as the sex of the connection. To tell if you have a null modem adapter or a gender changer, you would have to look for a label or check the pin configuration on either end of the adapter using a continuity tester or multi-meter.

Multimedia Connectors

Multimedia is one of the new hot areas of computers and how people are using them. Several standard types of connections are used for multimedia purposes. Table 1-4 lists some standard multimedia connectors and their uses.

Table 1-4		Types of Multimedia Connectors
<i>Role</i>	<i>Connectors used</i>	
Analog audio	The main analog audio connector used in your computer is the $f/8$ -inch connector, which is the standard headphone, speaker, and microphone jack for most sound cards. This form of connector has been around for years as an audio connector.	
MIDI	The Musical Instrument Digital Interface (MIDI) connector has been in use since the 1980s, but the standard for it was only finalized in the early 1990s. Most electronic keyboards and many other instruments have MIDI connections on them, which are usually 5-pin DIN connectors; while the connector on a computer is a DB-15 connector. MIDI connections allow for time synchronization, sequencing, recording, editing, and playback of MIDI-formatted data. The MIDI file format is extremely efficient.	
Analog video and cable	Computers make use of the same video transfer cables that are used in the home AV market. These include coaxial cable that can carry audio and video signals; RCA cables that can carry either an audio or video signal — but, if you have an RCA connector on your computer, it is likely a yellow video connector — and S-Video, which is used for higher quality video signals and uses a mini-DIN 4 connector. These cables are used to connect your computer to a variety of video devices to either capture or display data.	
Digital audio	The main digital audio connection for your computer is a S/PDIF (Sony Philips Digital Interface Format) connection. Many motherboards and soundcards support this interface for multimedia applications. The S/PDIF connection is often made using coaxial, RCA, or optical connectors and cables.	

Getting an A+

In this chapter, you examine the role of different types of connectors. In the process, you discover the differences between serial and parallel communication, as well as the differences between USB and IEEE-1394 (FireWire). You find out about the different types of cables and connectors that are used in different areas of a computer system. You should now have a general understanding about where and why different cables and connectors are used in your system.

Prep Test

- 1** Which of the following ports offers the fastest throughput or transfer speed?
 - A Serial
 - B Parallel
 - C USB
 - D IEEE-1284

- 2** What type of connector on the back of your computer accepts a parallel cable?
 - A DIN-8 female
 - B DB-25 female
 - C DB-9 female
 - D DB-25 male

- 3** What type of port uses a three-row DB-15 connector?
 - A Serial
 - B SVGA
 - C Parallel
 - D Game

- 4** RS-232 is a term associated with which type of port?
 - A Serial
 - B Parallel
 - C SCSI
 - D Game

- 5** Ethernet network cards typically use what type of connector?
 - A RJ-32
 - B DB-9
 - C DB-25
 - D RJ-45

- 6** BNC connectors are usually used on your computer to provide which of the following?
 - A Serial connections
 - B Parallel connections
 - C Networking
 - D Never used

7 How is serial data moved?

- A Across multiple cables at the same time
- B As a sequential stream of data
- C Using a bidirectional information algorithm
- D Using multi-clock sequencing

8 Data transfer speed for the IEEE-1394 standard is:

- A 10 Mbps
- B 100 Mbps
- C 400 Mbps
- D 12 Gbps

9 Most keyboards connect to computers by using which type of connector?

- A DB-9
- B DIN 6
- C Mini-DIN 6
- D RJ-45

10 Monitors usually connect to a computer through what type of connector?

- A DB-15
- B DB-25
- C USB
- D RS-232

11 CAT5 cables are defined by which of the following?

- A Number of pairs of wires
- B Length of the cable segment
- C Type of shielding that is used
- D Number of twists per foot in each pair of wires

12 Thin coaxial cable usually implements which type of connectors?

- A BNC
- B RJ-45
- C RS-232
- D IEEE-488

13 Pin 1 on a ribbon cable is usually identified by what?

- A The number 1 inscribed on the connector.
- B A colored line on the wire leading to it.
- C A keyed connector.
- D Pin 1 is always on the left side of the connectors.

Answers

- 1 C. USB offers the highest transfer rates. USB 1.1 has transfer rates of 12 Mbps, while USB 2.0 offers transfer rates of 480 Mbps. This question might have been disputable if IEEE-1394 (FireWire) had been in the question because it has transfer rates of 400 Mbps. *See "Universal Serial Bus (USB)."*
- 2 B. On the back of your computer, the parallel port has a DB-25 female connector. *Review "DB-25."*
- 3 B. VGA and SVGA use HD DB-15 connectors that have three rows of five pins. *Take a look at "DB-15."*
- 4 A. RS-232 is one of the standards associated with the serial port. *Peek at "Serial and parallel ports."*
- 5 D. RJ-45 is typically used for networking. *Look over "RJ-45."*
- 6 C. BNC connectors are used for Thinnet or 10Base2 Networking. *Study "BNC."*
- 7 B. Serial cables move data sequentially, one bit at a time, while parallel cables allow for all eight bits to be sent at one time. *Refer to "Serial and parallel ports."*
- 8 C. FireWire, or IEEE-1394, allows for transfer rates up to 400 Mbps. *Examine "FireWire (IEEE-1394)."*
- 9 C. Older PCs use DIN 5 connectors, while newer ones use the Mini-DIN 6 or PS/2 connector. *See "PS/2 or Mini-DIN 6."*
- 10 A. VGA and SVGA monitors usually connect to your computer via a high density DB-15 connector. *Review "DB-15."*
- 11 D. CAT5 cables are defined by the type of copper wire that is used and the number of twists per foot by each pair of wires. *Check out "Twisted pair."*
- 12 A. BNC connectors are used with Thinnet Ethernet networking. *Peruse "BNC."*
- 13 B. Pin 1 is identified by a colored line on the side of the cable that contains Pin 1. Keyed connectors prevent you from putting the cable in the wrong way, but do not specifically identify Pin 1. *Take a look at "Ribbon."*

Chapter 2: Installing and Configuring Input Devices

Exam Objectives

- ✓ Understanding keyboards
- ✓ Working with the mouse
- ✓ Modems and network cards
- ✓ Other input devices


The computer processes information, but only after someone inputs that information into the system. In this chapter, I discuss different types of input devices that are used to input data into the computer.

There are many forms of input devices that you can use to communicate with the computer, such as modems, network cards, touch pads, and biometric devices, but the keyboard and mouse are the most popular.

Minding Your Keys and Os

The keyboard is the main device you use to interact with a computer system. Your keystrokes are converted into the characters that you see on the screen and eventually print out. Most keyboards have five major key parts, or groups, that contain keys of a specific purpose. These areas are shown in Figure 2-1.

Figure 2-1:
The five major groupings on a keyboard.



The five major groupings on a keyboard are:

- ◆ **Alphanumeric:** The alphanumeric keys are the keys on the keyboard that contain the letters of the alphabet, numerals, and punctuation.
- ◆ **Function:** The 12 function keys on a keyboard offer special features. For example, in Windows, pressing the F2 keystroke is used to rename a file or icon.
- ◆ **Cursor:** The cursor keys allow you to move the cursor or insertion point around by using the keyboard.
- ◆ **Numeric:** The numeric keypad has the mathematical operators and numbers for quick, single-handed access for inputting numeric information.
- ◆ **LEDs:** The LEDs indicate whether features such as Caps Lock, Num Lock, and Scroll Lock are turned on or off.

Each key on the keyboard has a *keyswitch* that closes an electrical circuit on a grid when a key is pressed. When the key contacts the grid, the keyboard controller detects the keystroke and generates a *keycode*. The keycode is then converted into an ASCII code that is used to display the character on the screen.

Each location on the grid corresponds to a specific keycode. This is why you can't simply move the key caps around and expect to move the character that types. For example, the letter *A* is on the left side of the keyboard. Moving the *A* key to the right side of the keyboard will not allow you to press that key to type the letter *A* because the location on the grid that generates a keycode for letter *A* is on the left side of the keyboard.

The different keyboards work in different ways when you press a key. Each manufacturer decides what type of keyboard to manufacture and how the keystrokes will be interpreted by the system. The two most popular techniques used to identify what keys are being pressed are as follows:

- ◆ **Switch-based:** A switch-based keyboard uses micro-switches for each key. The downfall of a switch-key keyboard is that the keys deteriorate with time and tend to get dirty, but the good thing is that these keyboards are not expensive.
- ◆ **Capacitive:** A capacitive keyboard is also known as *membrane keyboard* and is more expensive than switch-based keyboards. Although a capacitive keyboard is more expensive, it is also more reliable. With a capacitive keyboard, each key pushes a spring, which pushes a paddle and creates an impression on the capacitive surface located under the keyboard. The impression on the capacitive surface sends a signal that is interpreted by the keyboard controller. This is the most common type of keyboard in portable computers.

In order for the keyboard to work with the system, there must be some software that drives the keyboard actions. Two types of software routines are used to allow the keyboard to work with the system:

- ◆ **Keyboard device driver:** Like any piece of hardware, there is a driver in Windows that is responsible for allowing the device to work in Windows — including the keyboard.
- ◆ **Firmware:** The keyboard firmware is typically stored in ROM in the actual keyboard (as in the case of the XT keyboard) or in a chip on the motherboard. The firmware contains low-level code to communicate with the hardware.

Identifying keyboard types

A number of different types of keyboards have been used over time. The list below identifies the different types of keyboards you need to know for the A+ Certification exam:

- ◆ **XT keyboard:** The older XT keyboard has 83 alphabetic keys which include 10 function keys. This keyboard had the keyboard processor located on the keyboard itself. In addition to the alphabetic keys, it had a numeric keypad and cursor control keys on the right side of the keyboard.
- ◆ **AT keyboard:** The older AT keyboard has 84 regular keys (an extra “System Request” key was added), which again includes 10 function keys. The AT keyboard has a bigger Enter key than the XT keyboard, and the keyboard processor moved from the keyboard itself to the computer’s motherboard. The AT keyboard uses the *AT keyboard connector*, which is also known as the *DIN-5 connector*.
- ◆ **Enhanced keyboard:** The enhanced keyboard is the popular keyboard used today and has 101 keys and includes 12 function keys. The enhanced keyboard has a numeric keypad along with cursor control keys located on the right side. The enhanced keyboard typically uses the *PS/2 connection* (also known as a *Mini-DIN 6 connector*) or a USB connection.
- ◆ **Windows keyboard:** A Windows keyboard is similar to an enhanced keyboard containing 101 keys, but the Windows keyboard contains buttons, or keys, that control features of Windows. For example, on a Windows keyboard, there is a button to pop up the Start menu and a key to pop up the shortcut menu, as if you had right-clicked with the mouse.
- ◆ **Natural or ergonomic keyboards:** The *natural*, or *ergonomic*, keyboards are a modification of the enhanced keyboards. The natural keyboard separates the alphabetic keys into two parts: one part for the left hand and the other for the right hand. The keyboard is bowed and may sometimes come apart to allow for natural placement of the hands. These keyboards normally also have a wrist rest to help your hands sit in a more natural position. These keyboards are designed to reduce repetitive stress injuries. Figure 2-2 shows a typical natural keyboard.




Figure 2-2:
A natural,
ergonomic-
style
keyboard.

Installing a keyboard

Installing a keyboard is fairly straightforward — you pretty much plug the keyboard into the system, and Windows uses a keyboard driver to communicate with the device.



For the A+ exam, it is important to know the different keyboard connectors, which are described in the following bulleted list:

- ◆ **DIN-5 connector:** The *DIN connector* is also known as the *AT connector* and is used on older systems to connect the keyboard to the computer. You can always tell the DIN connector because it is large and round.
- ◆ **PS/2 connector:** The *PS/2 connector*, also known as the *Mini-DIN 6 connector*, is the most popular keyboard connector on today's computers. It is identical to what is found on a typical mouse. Figure 2-3 shows a DIN-5 connector beside a PS/2 connector.
- ◆ **USB connector:** Keyboards that connect to the system via a USB connector are becoming very popular. A number of systems today are getting away from the PS/2 connections for keyboards and mice and using USB. Figure 2-4 shows a keyboard being connected to a USB port.

Configuring keyboards

Configuring a keyboard in Windows doesn't involve a lot of options — normally you just plug it in and it works. Windows does allow you to configure the blink rate of the cursor and lets you change the repeat rate, which dictates how quickly a letter will repeat if you hold a key down on your keyboard.

Figure 2-3:
A DIN-5
connector
(left) beside
a PS/2
keyboard
connector
(right).



Figure 2-4:
A USB
keyboard
connects to
a USB port
on the
system.

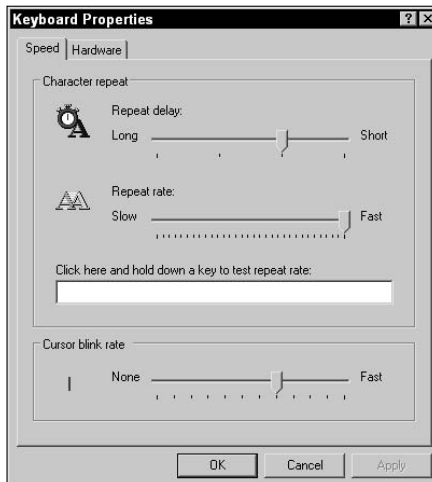


To configure the keyboard options in Windows XP, follow these steps:

1. Choose **Start**⇨**Control Panel**⇨**Printer and Other Hardware**⇨**Keyboard**.

The Keyboard Properties dialog box appears, as shown in Figure 2-5.

Figure 2-5: Changing keyboard properties through the Control Panel in Windows XP.



2. Adjust the settings to suit your needs.

Here you can configure the character repeat options and the cursor blink rate. The following options are available:

- **Repeat Delay:** The repeat delay controls the length of time between when you first hold down a key and when the character begins to repeat on-screen.
- **Repeat Rate:** The repeat rate controls how quickly additional characters are displayed when the key is held down.
- **Cursor Blink Rate:** The cursor blink rate controls how fast the cursor blinks.

3. Click OK to apply your settings and close the dialog box.

There are also a number of accessibility features in Windows that deal with the keyboard. Accessibility features in Windows are designed for users who have disabilities. To configure the keyboard accessibility options in Windows XP, follow these steps:

1. Click **Start**→**Control Panel**→**Accessibility Options**.

The Accessibility Options dialog box appears, as shown in Figure 2-6.

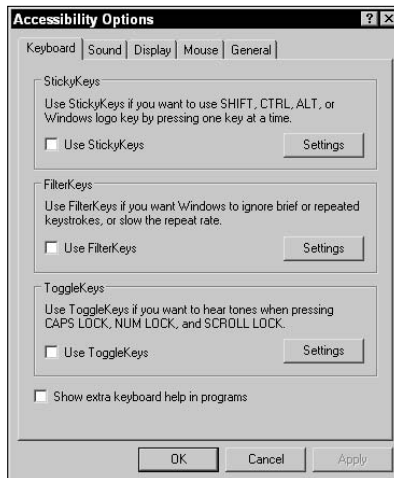
2. Ensure that the **Keyboard** tab is selected and then configure the following options as needed:

- **StickyKeys:** When StickyKeys is enabled, when you press the Shift, Ctrl, Alt, or (if you have a Windows keyboard) Windows logo key, the key remains active, as if you were holding it down, until you press it again. This allows someone who cannot press two keys at the same time to still take advantage of these keys.

- **FilterKeys:** Use FilterKeys to tell Windows to ignore repeated keystrokes. For example, if you hold down a key too long with FilterKeys enabled, the repeat rate of the keystroke is reduced. This will prevent the character from appearing many times at a quick rate.
- **ToggleKeys:** ToggleKeys is an option that triggers audible tones when the Caps Lock, Num Lock, or Scroll Lock keys are pressed.

3. Click OK.

Figure 2-6: Changing the keyboard accessibility options in Windows XP.



Catching the Mouse

The other popular input device found on computer systems today is the mouse. The mouse is one of the primary input devices due to the *graphical user interface (GUI)* features of today's operating systems. The following sections identify the different types of mice and how to install them.

Types of mice


You can find three types of mice on today's desktop computer systems: mechanical mice, opto-mechanical mice, and optical mice.

- ◆ **Mechanical:** A mechanical mouse was the only type of mouse used with computers for many years. The mechanical mouse uses a rubber ball that moves a pair of wheels inside the mouse that control, or capture, the horizontal and vertical movements of the mouse and then send that information to the computer. The movement is then displayed on the screen.
- ◆ **Opto-mechanical:** The opto-mechanical mouse is a very popular mouse today. It uses a rubber ball that moves rollers that control the movement

of a disc, known as the encoding disc, which has holes along the side of it. The term “opto” comes from the fact that there are LEDs and sensors that use an infrared light beam to determine the movement of the mouse.

As shown in Figure 2-7, the sensors in this type of mouse use infrared beams. When the rubber ball moves the encoding disc, the infrared beam passes through the holes on the encoding disc indicating that the mouse has moved and in which direction.

Figure 2-7: As the mouse moves, the holes and solid areas of the encoding disc help determine the movement of the mouse.



- ◆ **Optical:** An optical mouse is different than an opto-mechanical mouse because the optical mouse has no moving parts. The rubber ball and wheels are replaced with an optical scanning system that works on pretty much any surface.

The optical scanning system is made up of optical sensors and a *Digital Signal Processor* (DSP). The optical mouse captures images of the surface at a rate of 1500 images per second. The DSP then analyzes and compares the images to detect the mouse movement.

Installing a mouse

Installing a mouse is very similar and just as easy as installing a keyboard. You pretty much just plug the mouse in and the Windows mouse driver controls the device. The following is a list of the different types of interfaces you may use to connect a mouse to the system.

- ◆ **Serial:** A serial mouse is a mouse that connects to the serial port on the computer. The serial port is the male 9-pin or 25-pin port on the back of the computer that sends data one bit at a time. Serial mice are not as popular today as they once were.

- ◆ **PS/2:** One of the most popular interfaces for a mouse on desktop computers is the PS/2 connection. The PS/2 connection is also known as the Mini-DIN 6 connector.
- ◆ **USB:** Becoming the most popular type of connection for mice today is the USB connector. Simply plug the mouse into the USB port and click away! As you find out in many chapters of this book, USB is hot-swappable — meaning that you can plug and unplug a USB device without shutting down the system. USB can transfer data at 12 Mbps (USB 1.1) or 480 Mbps (USB 2.0).

Communicating with Modems and Network Adapters

In this section, you find out about installing modems and network cards that allow computer systems to communicate with other systems on the network.

Working with modems

A *modem* is a communication device that *modulates*, or converts, the digital signal from the computer into an analog signal so that it can travel over an analog line, such as telephone wire. The signal is then *demodulated* back into a digital signal at the modem on the receiving system. The process of *modulating* and *demodulating* is where the word *modem* comes from.

The modem allows a computer to send data over an analog line and is useful when you want to dial up another computer over the phone line or dial into the Internet.

Modem characteristics

A modem is either internal or external:

- ◆ **An internal modem** is installed in an expansion slot within the computer housing as either an older ISA card or, more popular today, a PCI card.
- ◆ **An external modem** sits on the desk beside the computer and plugs into the computer's serial port.

Both internal and external modems plug into a wall telephone jack via a standard telephone cable.

Another very important characteristic of the modem is its speed. The speed of the modem is measured in *bits per second (bps)* — and like anything else, the more bits the better! When it comes to modem speed, you're looking at 56 Kbps (56,000 bps) as the standard. This is known as a 56K modem.

AT commands

For the A+ exam, you need to be familiar with what an AT command is. AT is short for *Attention*, and AT commands are the modem commands that are called by software to perform communication. You typically don't need to use these commands unless you're troubleshooting why a computer cannot dial up. Your first question is always, "Is it the modem or the application I am using?" By using these low-level AT commands, you can determine whether the modem is working.



The AT commands that you should be comfortable with for the A+ exam are listed in Table 2-1.

Table 2-1 **AT Modem Commands**

<i>Command</i>	<i>Description</i>
ATDT <number>	Dial a number using touch-tone dialing.
ATDP <number>	Dial a number using pulse or rotary dialing.
ATA	Instructs the modem to answer.
ATH	Instructs the modem to hang up.
ATL	Sets the speaker loudness.
ATZ	Resets the modem to the default settings.

There are a few tricks to remembering these commands for the A+ exam. Notice that each command starts with the *AT* prefix, and then you have a character such as "D" for dial, "A" for answer, or "H" for hang-up. Because there are two types of dial tones, the dial command has another character representing the type of dial tone to use and requires the telephone number to dial. For example, to call Ed's (the other author of this book) house with touch-tone dialing, the AT command would be

```
ATDT 555-5555
```

Okay, this isn't Ed's phone number. He wouldn't like it too much if I put his personal phone number in the book. But you get the idea of how the command works. There are more AT commands than just what I show here, but these are the ones you need to be familiar with for the exam.

Installing a modem

Most modems today ship as PCI cards and are inserted into a PCI slot in the computer — this is known as an internal modem. You can install an internal modem by following these steps:

- 1. Power off the computer and be sure to ground yourself.**
- 2. Remove the cover from the computer.**

- 3. If necessary, remove the blanking plate that allows access to the modem from the back of the PC.**

This allows the modem card to fit in the slot.

- 4. Place the modem card into the appropriate expansion slot.**

For example, you will most likely have a PCI card, so place the PCI modem card into an available PCI slot.

- 5. Screw the modem card into place and put the cover back on the PC.**

- 6. The modem has two ports on it that both take RJ-11 connections — connect the incoming phone line to the appropriate port and then connect the phone to the modem's second port.**

- 7. Power on the PC.**

Configuring a modem

After you have installed the modem and Windows starts up, Plug and Play should kick in and detect the new device. If a driver is not present in Windows for the device, you will be asked for the driver. If you don't have the driver, you can choose to cancel adding the driver at this time and go to Device Manager at a later time to update the driver after you have it. See Book VI, Chapter 1, to find out how to load and configure drivers by using Device Manager.

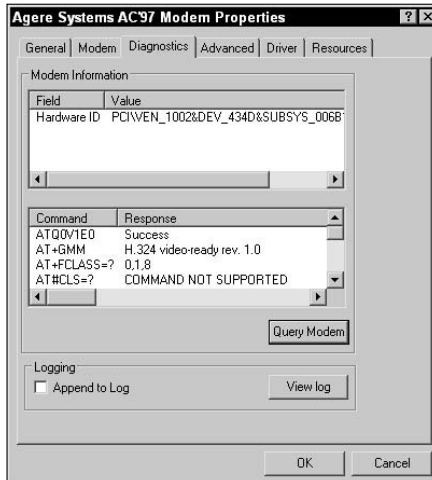
When the driver has been loaded, you can go to the properties of the device in Device Manager and query the modem to verify that the modem is functioning properly. *Querying the modem* is the term used to describe sending AT commands to the modem and verifying that the modem responds to the commands. The nice thing is you don't need to type each command listed earlier in Table 2-1 — there is a Query button in the properties of the modem that does this for you. To verify that the modem is working, follow these troubleshooting steps:

- 1. Click Start and right-click My Computer.**
- 2. Choose Properties from the pop-up menu.**
- 3. In the System Properties window, click the Hardware page tab and then click the Device Manager button.**
- 4. In Device Manager, expand the modems node, right-click your modem, and choose Properties.**
- 5. In the Properties window of the modem, click the Diagnostics page tab (shown in Figure 2-8) and then click the Query Modem button to perform a diagnostic on your modem using the AT commands.**

You should notice a list of AT commands that have been issued against the modem as a diagnostic test and whether the test was successful or not.

- 6. Close all windows.**

Figure 2-8:
Performing
trouble-
shooting
diagnostics
on the
modem in
Windows
XP.



Working with network adapters

The most popular communication device found in computers today is the network card. A network card is a device that allows you to connect your computer to a network or the Internet, and is responsible for sending and receiving data on the network. Not only is the network card popular in systems found in the corporate world, but because of the popularity of high-speed Internet, you can also find them in a lot of home computers as well.

Network cards are categorized by their transfer rates. For example, a network card that can transfer 10 million bits of data per second is known as a 10 Mbps network card, while a network card that can transfer 100 million bits per second is known as a 100 Mbps network card.



It is important to note that the card will run at the highest common speed available between the two communicating devices. For example, if you connect your system to a 10 Mbps network hub, and your system has a 100 Mbps network card, then your card will run at only 10 Mbps. This is why most systems have what is called a 10/100 network card — the card can talk at either 10 Mbps or 100 Mbps.

Installing a network adapter (card)

When installing a network card, you first need to determine what type of interface you wish to use. There are a number of types of network cards that could be installed into the system: you could install a USB network card, an old ISA card for older computers, or more popular today is the PCI network card (shown in Figure 2-9) that is inserted into a PCI slot.




Figure 2-9: Installing a PCI network card into the PCI slot of a computer.

To install a network adapter, follow these steps:

- 1. Power off the computer and be sure to ground yourself.**
- 2. Remove the cover from the computer.**
- 3. If necessary, remove the blanking plate that allows access to the modem from the back of the PC.**
This allows the network card to fit in the slot.
- 4. Place the network card into the appropriate expansion slot.**
For example, you will most likely have a PCI card, so place the PCI network card into an available PCI slot.
- 5. Screw the network card into place and put the cover back on the PC.**
- 6. Connect the network cable to the back of the network card.**
- 7. Power on the PC.**

Configuring a network card

After you install the network card, Plug and Play should kick in and either load the driver for you or prompt you for the driver. When the driver is loaded, you are able to specify settings for the network card in Windows. For example, if

you want to force your network card to run at a particular speed, you can do that through the properties of the network card.

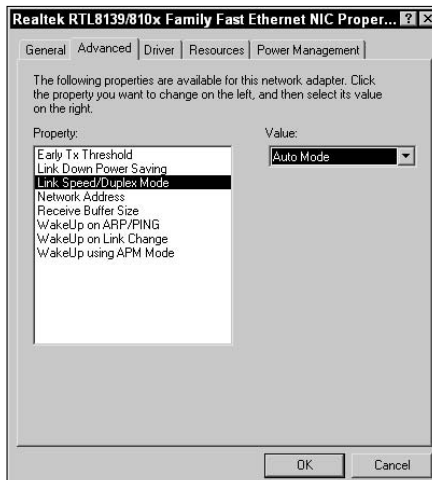
To configure your network card, follow these steps:

1. Click Start and right-click My Computer.
2. Choose Properties from the pop-up menu.
3. In the System Properties window, click the Hardware page tab and then click the Device Manager button.
4. In Device Manager, expand the Network Adapters category on the left, right-click your network card, and choose Properties.
5. In the Properties dialog box for the network card, click the Advanced page tab and you will see settings that can be configured for the network card (shown in Figure 2-10).

The advanced settings will be different for each make and model of network card because the settings are particular to the driver that is loaded. A very popular setting is the link speed, which specifies what speed you wish to run the card at. Most network cards are set to Auto, which means the card will detect the best speed — but when troubleshooting, you can force a particular speed.

Typically, you may also choose which type of connector on the network card will be used, if it has multiple connector styles. For example, a number of network cards ship with both an RJ-45 port and a BNC port — this is called a *combo card*. If you want to ensure that the system is using the RJ-45 port, then you set that in the device properties.

Figure 2-10:
Configuring
network
card
settings
through
Windows
Device
Manager.



Other Input Devices

As technology moves on, new input devices appear. For example, laptops use a touch pad as an input device, whereas handheld devices use a pen-like stylus. The following is a list of other input devices you may encounter:

- ◆ **Touch pad:** A number of laptop systems today use the touch pad as an alternative way to control the mouse pointer in Windows. The touch pad is a rectangular area found below the keyboard that you move your finger across in the direction you would like the mouse pointer to move. There are usually two buttons below the square to perform left- and right-click operations. A touch pad is displayed in Figure 2-11.
- ◆ **Touch screens:** You see touch screens used a lot at public terminals for things like information booths and bank machines. You can identify the touch screen by it being a device that allows you to actually *touch the screen* to choose different options.
- ◆ **Bar code reader:** A bar code reader is a type of scanner that reads bar codes and then converts the bar code into data that is stored on the computer and used in applications. An example of an application that may take advantage of a bar code reader as an input device is an inventory application.



Figure 2-11: A touch pad is a popular input device to control the mouse pointer on a laptop computer.

- ◆ **Biometric devices:** Biometric devices are used to authenticate, or log, people onto a system by using biological characteristics of that person, such as a fingerprint or a retinal scan. Biometric devices are very popular in high-security environments because of how difficult it is to impersonate someone when dealing with authentication methods such as a fingerprint scan. Some laptop makers are even starting to include fingerprint scanners in consumer-level laptops as security devices and theft deterrents.

Getting an A+

In this chapter, you find out about different types of input devices. The following list outlines some of the key points to remember for the exam:

- ◆ Keyboards usually connect to the system with a PS/2 connection or a USB connection with today's systems.
- ◆ Older keyboards use the DIN-5 to connect to the motherboard.
- ◆ Most mice today connect to the system via a USB port or a PS/2 port.
- ◆ A serial mouse connects to the serial port, but this type of mouse isn't very common anymore.
- ◆ Modems are used to convert digital signals to analog and vice versa so that digital data can be delivered over an analog line.
- ◆ Network cards are the most popular type of communication device in today's systems. Be sure to review the networking chapters to fully understand how to network a system.

Prep Test

- 1** Which two formats do serial ports come in?
 - A 9-pin
 - B 15-pin
 - C 20-pin
 - D 25-pin

- 2** At what speed does USB 1.1 transfer information?
 - A 10 Mbps
 - B 12 Mbps
 - C 100 Mbps
 - D 480 Mbps

- 3** Which type of keyboard is designed to reduce repetitive stress injuries?
 - A AT keyboard
 - B Natural keyboard
 - C Windows keyboard
 - D Enhanced keyboard

- 4** Which type of mouse uses a rubber ball and sensors connected to two wheels to determine mouse movement?
 - A Mechanical
 - B Opto-mechanical
 - C Optical
 - D Wheel

- 5** What is the AT command to hang up the connection?
 - A ATDT
 - B ATA
 - C ATH
 - D ATDP

- 6** At what speed does USB 2.0 transfer information?
 - A 10 Mbps
 - B 12 Mbps
 - C 100 Mbps
 - D 480 Mbps

Installing and Configuring Input Devices

7 Which two popular connector types are used with a mouse?

- A USB
- B Parallel
- C PS/2
- D OS/2

8 Which two types of interfaces are typically used to connect a network card to the system?

- A Parallel
- B USB
- C PCI
- D Serial

9 What is the AT command to call 555-5555?

- A ATDT 555-5555
- B ATA 555-5555
- C ATH 555-5555
- D ATZ 555-5555

10 What type of input device is used by a warehousing application to take inventory of all the products?

- A Touch pad
- B Touch screen
- C Touch mouse
- D Bar code reader

Answers

- 1** A, D. Serial ports, which are used by serial mice, come in two formats: 9-pin and 25-pin. *See "Installing a mouse."*
- 2** B. USB 1.1 devices run at 12 Mbps, while USB 2.0 devices run at 480 Mbps. *Review "Installing a mouse."*
- 3** B. Natural keyboards are a form of ergonomic keyboard that reduces repetitive stress injuries. *Check out "Identifying keyboard types."*
- 4** A. A mechanical mouse uses all moving parts, such as a rubber ball that moves wheels that are picked up by sensors to detect mouse movements. *Peruse "Types of mice."*
- 5** C. The AT command used by modems to hang up a connection is ATH — the H stands for *hang up*. *Take a look at "AT commands."*
- 6** D. USB 2.0 runs at 480 Mbps, while USB 1.1 runs at a lower speed of 12 Mbps. *Peek at "Installing a mouse."*
- 7** A, C. Popular connections used by the mouse today are the USB connector and the PS/2 connector. *Look over "Installing a mouse."*
- 8** B, C. Network cards connect to the system as either a USB or PCI device. *Study "Installing a network adapter (card)."*
- 9** A. The ATDT command is used to dial a number. *Refer to "AT commands."*
- 10** D. A bar code reader is a type of device that reads bar codes and converts the code into data to be used on the computer. *Examine "Other Input Devices."*

Chapter 3: Installing and Configuring Output Devices

Exam Objectives

- ✓ Understanding video adapters and monitors
- ✓ Working with sound cards
- ✓ Other output devices

An *output device* is a device that sends data out of the computer and creates a tangible output. Examples of output for the system are displaying information on the monitor, printing out a report, or even listening to the sound that the sound card puts through the speakers when playing your favorite computer games.

This chapter also looks at the video hardware that is responsible for creating the images on the screen and discusses how a sound card works.

Understanding Video Adapters

The *video adapter*, also known as the *video card*, is the interface between the monitor and the computer and is responsible for converting the digital data from the computer into analog information. The data is converted to an analog signal before being delivered to the monitor because the monitor uses analog data to create the image.

Because display functions are very time- and memory-consuming, most video adapters these days have their own processing chip and memory to alleviate the processing workload from the CPU. Today's video adapters are typically AGP cards, but you may run into some older systems that require a PCI video card because these systems don't have an AGP slot.

The following outlines the basic role of what the video adapter does when it comes to displaying data on the computer screen:

1. Data is sent to the video card via the expansion bus that the video card resides in.
2. The video chipset on the video card writes the data to memory located on the card.

3. After being stored in memory, the data is passed to the *digital-to-analog converter (DAC)*, where it is converted from digital signals to analog signals that the monitor can understand.
4. The data is passed to the monitor, which then displays the data on the screen.

The video adapter is identifiable by its unique 15-pin female connector made up of 5 pins in three rows. On most systems today, the port on the video card is typically blue and has a monitor symbol beside it. Figure 3-1 shows a video port.

Figure 3-1: Looking at the back of a video card and the port style.



Looking at the video standards

A number of video standards have been developed over the years, with each standard increasing the quality of the display from the previous standards.

The following list outlines some of the popular video standards, and you will need to know them for the A+ Certification exam:



- ◆ **Monochrome Display Adapter (MDA):** The MDA standard displays data in a text format with a single color — white text on the black screen. This standard is pretty much obsolete, but you may encounter an old server that still uses it.
- ◆ **Color Graphics Adapter (CGA):** The CGA standard is the next step above MDA and supports four colors in a 320 x 200 resolution but supports only two colors with a resolution of 640 x 200. CGA video adapters were the first adapters to support color.
- ◆ **Enhanced Graphics Adapter (EGA):** The next step after CGA is EGA, which is much improved over the CGA graphics standard. EGA supports 16 colors at 640 x 350 resolution.

- ◆ **Video Graphics Array (VGA):** VGA is a video graphics standard that allows for a resolution of 640 x 480 with 16 colors, but also supports 256 colors at lower resolutions.
- ◆ **Super VGA (SVGA):** Super VGA supports 16 million colors at resolutions as high as 1280 x 1024 and is the graphics standard of today.

Video board features

Although there are a number of different video standards available and a wealth of different video cards to pick from, some features are common to every video card. This section outlines some important characteristics of video cards, such as resolution, colors, and video memory.

Built-in coprocessors

A number of new video cards have a *built-in coprocessor* that processes most of the display functions on behalf of the CPU. This allows the CPU to service other requests on the system and not have to worry about processing any video commands. These advanced video cards are also known as *accelerated video cards* or *graphics accelerator cards*.

Resolution

Resolution is a term that techies use to describe how functional a video card is. For example, I may brag up my video card as being able to run at 1024 x 768, and you could turn around and say, “Is that all? My video card supports 1280 x 768.”

So what does the resolution describe? Resolution describes how many pixels are supported by both the video adapter and the monitor. A *pixel* is a small dot that is joined with other small dots to help create the image on the screen. The more pixels that are supported on your screen, the closer the dots must be — which gives you a finer display.

As an example, a video card that supports 1280 pixels going across the screen and 768 pixels going down the screen (1280 x 768) will create a much finer image than a video card that supports only 800 pixels going across the screen and 600 pixels going down the screen (800 x 600).

Colors supported

The picture on your monitor will look much better if the video card supports a lot of colors, so one of the other often-touted video card characteristics is how many colors the video card supports. Bottom line, the more colors the better. If you're curious to see what the big deal is about the number of colors supported, here's a quick way to get a look at the difference:

1. Right-click a blank space on your Desktop and select Properties.

2. On the Desktop tab, select “Bliss” from the Background list.

Bliss is the default background that XP uses — the one with the clouds that looks like Teletubby land.

3. Click Apply to change to the Bliss background.

4. Click the Settings tab.

5. In the Color Quality drop-down list, choose 256 colors.

You may not be able to choose 256 colors because the driver may not let you change your color settings that low.

6. Click Apply.

You should notice that the wallpaper doesn't look good anymore because Windows cannot display all of the colors that appear in the image.

7. Change your Color Quality back to its original value.

8. Choose OK.

Video memory

Because a video card has to process such a large amount of data, the video card comes with its own memory — known as *video memory*. Video memory is available in many different sizes, and popular amounts are 64MB, 128MB, and 256MB. Some video cards simply use DRAM, but because the demands of video display are so high and DRAM was so slow in the past, other memory types were developed for video:

- ◆ **Video RAM (VRAM):** VRAM has the benefit of being *dual-ported* — it can be read from and written to simultaneously — which adds to the performance of the memory.
- ◆ **Multibank DRAM (MDRAM):** MDRAM is a type of video RAM that uses the full width of the video bus (expansion bus that the video card resides in) with fewer memory chips.
- ◆ **Windows RAM (WRAM):** WRAM stores its data in chunks, which makes the data transfer faster than the more-expensive VRAM. Like VRAM, WRAM is dual-ported.
- ◆ **Synchronous Graphics RAM (SGRAM):** SGRAM is a type of video memory that is *single-ported*, which means that you can only read from or only write to the memory at one time. SGRAM is four times faster than other types of DRAM memory and is synchronized with the CPU clock.

Being able to change some of the video card's characteristics, such as color support or resolution, is a characteristic of the video driver that is installed. You may find that you cannot set your video card to 24-bit color because the driver does not support it. Of course, the driver does not support it because the card does not support it. Before I show you how to change your resolution, you need to know how to install a video card, which is covered in the next section.

Installing a video card

Installing a video card is like installing any other card, such as a network card or modem. You simply need to place the card into a supported bus architecture, typically either PCI or AGP, and then load the driver.

The following steps outline how you install a video card:

- 1. Power off and unplug the computer.**
- 2. Open the case or side panel of the computer so that you can access the expansion slots.**
- 3. If you are replacing a video card, take out the old one.**

If the old video card is integrated into the motherboard, you will not be able to remove it, but you may be able to disable it in CMOS. (For more information on CMOS check out Book II, Chapter 4.)

- 4. If you are adding the first video card to the system, then remove the blanking plate for the AGP or PCI slot you wish to use.**
- 5. Insert the card into the slot until it fits firmly.**
- 6. Screw the video card's plate into place to secure the card.**
- 7. Place the cover back on the computer, plug it back in, and power it on.**

When you power up the system, Plug and Play should kick in and detect the newly added device and either load a driver or prompt you for the driver. After the driver is loaded, you will be able to configure the display settings that are shown in the following sections.

Using Your Monitor

After you install the video adapter (see the previous section), you then need to connect the monitor to the video adapter. The following sections identify the different types of displays that are used with computers today and then discuss how to configure your display through Windows.

Types of displays

There are a few different types of displays that you may find connected to a system: you may encounter a CRT monitor, an LCD monitor, or even a projector. In this chapter, I focus on CRTs because LCDs are discussed in Book III, Chapter 7, and you are not likely to get any questions on the A+ exam about projectors. I do define each type of display, though.

Cathode Ray Tube (CRT)

Cathode Ray Tube (CRT) monitors were popular for computing for many years and were also popular as television screens. After receiving a signal from the video card, a CRT monitor displays the image by using an electron gun that shoots electrons at the phosphors covering the back of the screen, causing those areas of the screen to glow.

The electron gun located in the CRT continues to fire at the back of the screen from left to right and from top to bottom. This causes a glow of three phosphors colored red, green, and blue, to create a single pixel on the screen. The combination of all the lit pixels creates the image that you see, and the speed at which the pixels change gives the illusion of moving objects on the screen.

You may have heard the terms *refresh rate* and *dot pitch*, two very important terms used to describe a monitor. A monitor's *refresh rate* refers to how often the electron gun can redraw the entire screen — the faster it can redraw the screen, the smoother any moving objects appear.

A monitor's *dot pitch* is simply the distance (in millimeters) between two pixels of the same color on a monitor. For example, an average monitor has a dot pitch of 0.28 millimeters, while a better monitor has a dot pitch as tight as 0.24mm.

LCD

A *Liquid Crystal Display (LCD)* is a flat monitor type that has been used with laptops for many years and is now used for flat-screen desktop monitors. You can find out more about LCDs in Book III, Chapter 7.

Projector

I can't really talk about display types without talking about using a projector to display the image. Screen projectors are becoming popular both for home theater systems and in boardrooms.

When using a computer or laptop with a projector, the projector connects to the video adapter via a 15-pin video port. If you're using a laptop, you may notice that the output does not come out of the projector right away. Most laptops allow you to choose whether to display the image on the LCD, projector, or both by holding down the function (FN) key and then pressing either F4 or F5. Look for the key that has a monitor on it to get your projection going.

Configuring your display settings

You can change the characteristics of your display — such as color support or resolution — after the driver has been loaded for your video adapter. Whatever settings you choose must be supported by both the video card and the monitor. For example, choosing too high a resolution can distort the display. In this example, the resolution is supported by the video card but not the monitor. Mismatched resolution was such a problem in the past that Windows operating systems now ask you whether the display is okay when you change the resolution. If you choose Yes, then the settings are applied — if you choose No or don't choose anything for a few seconds, the settings are switched back.

To configure your display settings in Windows XP, follow these steps:

1. Choose Start⇨Control Panel⇨Appearance and Themes⇨Display.
2. Click the Settings tab (shown in Figure 3-2).

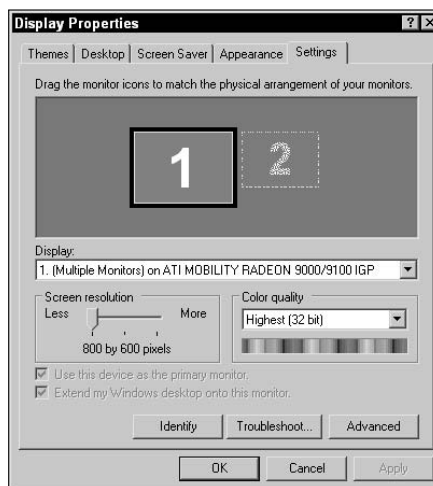


Figure 3-2:
Configuring the display settings in Windows XP.

- 3. From the Color Quality drop-down list, choose how many colors you want your display to support.**
 - *16-bit* supports 65,536 colors.
 - *24-bit* supports 16 million colors.
 - *32-bit* supports over 16 million colors.
- 4. In the Screen Resolution section, move the slider to set the resolution of your monitor to the desired setting.**
- 5. Click OK.**
- 6. When asked if the display looks okay, click Yes.**

Configuring multi-display support

Windows 2000/XP and Windows Server 2003 support a feature called multi-display support. Multi-display support allows you to insert multiple video cards into the system and have a monitor connected to each video card. Once a monitor is connected to each video card, you can then *extend* your Windows desktop to include the second monitor. This will allow you to use the screen space from both monitors at the same time.

Multi-display support allows you to do things such as move one program to one of the monitors and then use another program on the second monitor. For example, I place my e-mail program on my right-hand monitor and surf the Internet on my left-hand monitor. When an e-mail comes in, I already have the e-mail program open so I can simply read the e-mail.

To configure multi-display support in Windows, follow these steps:

- 1. Open your computer case and install an additional video adapter.**

You can find instructions for installing a video card in the “Installing a video card” section.
- 2. Put the computer case back on and power on the system.**
- 3. Plug a separate monitor into each of the two video cards.**
- 4. Ensure that a driver is loaded for each video card and for each monitor.**
- 5. To enable multi-display support, right-click on the Desktop and choose Properties.**
- 6. In the Desktop Properties window, click the Advanced page tab and select the second monitor (contains the number two in the icon).**

7. Once you have highlighted the second monitor, select **Extend My Windows Desktop to This Monitor**.
8. Click **OK**.
9. Start **Internet Explorer**, type `www.gleneclarke.com` in the address bar, and press **Enter**.
10. Once the Web site appears, restore down the window and move it to the second monitor.
11. Start your e-mail program and ensure it is on the first monitor.



Lab 3-1 demonstrates how to configure Windows XP for multi-display support after installing two video cards in the system. Lab 3-1 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

The Sound of Computers

Another very popular output device is the sound card. As an output device, the *sound card* is responsible for sending sound data to the speakers, but can also be used as an input device by lining-in an external source to the card. For example, you could line-in a stereo system through your sound card and convert some old audio tapes to MP3 files.


The sound card contains a *digital-to-analog converter (DAC)* that converts the digital data from the computer into analog signals that create the sound. The sound card has a number of ports on the back that allow you to connect both output devices and input devices:

- ◆ **Line-in:** The Line-in port is used to connect your CD player or stereo system and use it as a source for what the sound card plays. A great idea to get your old tapes to MP3 format!
- ◆ **Mic:** Used to connect a microphone, which acts as a source for information being recorded.
- ◆ **Headphones:** Used to connect your headphones.
- ◆ **Line-out:** Just as you can connect your stereo to the computer via the Line-in port, you can also take sound on the computer and “line it out” to the stereo.

Types of sound cards

You can find a number of different types of sound cards. The most popular interfaces for a sound card are as follows:

- ◆ **Industry Standard Architecture (ISA):** ISA cards were a popular choice on older systems and in the early Pentium days, but you still may encounter one from time to time. An ISA card runs at 8 MHz and is only 16-bit.
- ◆ **Peripheral Component Interconnect (PCI):** Today's systems use PCI sound cards that are Plug and Play and are inserted into the PCI slots — those are the white slots on the motherboard. PCI devices outperform ISA by running at 33 MHz and are 32-bit cards.
- ◆ **Integrated:** Some motherboards also have *integrated* sound cards — meaning that the motherboard has the sound card built-in and you don't need to purchase an additional card to install in the computer.



A classic A+ Certification exam question asks what the default IRQ (interrupt request) is for sound cards — the answer is IRQ 5. You can find more information about IRQs and their roles in Book III, Chapter 4.

Installing a sound card

Installing a sound card is no different than installing any other type of card. You first need to decide what type of card you want to use in your system (and then go buy one). With sound cards, you are most likely using a PCI card, but you may encounter older ISA sound cards if you're supporting older systems.

Here's how you install a sound card:

- 1. Power off and unplug the computer.**
- 2. Open the case or side panel of the computer so that you can access the expansion slots.**
- 3. If you're replacing a sound card, take out the old one.**

If the old sound card is integrated into the motherboard, you won't be able to remove it and you may have to disable it in CMOS.
- 4. If this is the first sound card for the system, remove the blanking plate for the ISA or PCI slot you wish to use.**
- 5. Insert the card into the slot until it fits firmly.**
- 6. Screw the sound card's plate into place to secure the card.**
- 7. Place the cover back on the computer and power up the computer.**

Plug and Play should detect the device. After Plug and Play detects the sound card, it may prompt you for the driver — if prompted for the driver, supply the CD that came with the card. If you aren't prompted for the driver and Windows did not load a driver, then you may need to run the driver setup program from the CD that came with the card.

Configuring a sound card

After the driver has been loaded for your sound card, you can configure the sound card by going to the Sound applet in the Windows Control Panel. To configure the sound card, choose Start⇨Control Panel⇨Sound, Speech, and Audio Devices, and then choose Sounds and Audio Devices to bring up the Sound and Audio Devices Properties dialog box, as shown in Figure 3-3.

The following list summarizes each of the page tabs and the types of settings that you may configure in the Sound and Audio Devices Properties dialog box:

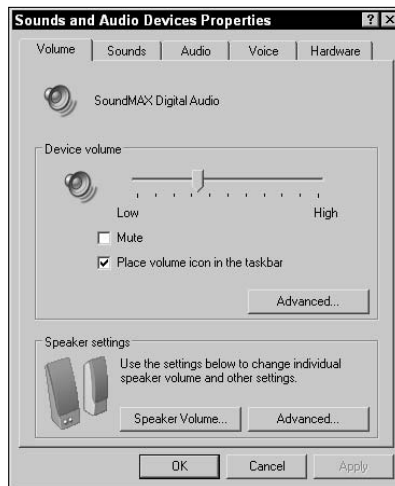


Figure 3-3: Configuring sound settings in Windows XP.

- ◆ **Volume:** Allows you to set the master volume for the sound card and choose options like muting the sound card. You may also choose to have an icon display in the taskbar — this allows for quick access to the sound card properties.
- ◆ **Sounds:** On the Sounds tab, you can specify your sound theme, which will set different sounds for different events in Windows, such as when Windows starts or shuts down.
- ◆ **Audio:** On the Audio tab, you can specify what sound device to use for playback or recording.
- ◆ **Voice:** On the Voice tab, you can specify what device to use when capturing or playing back voice recordings.
- ◆ **Hardware:** On the Hardware tab, you can get quick access to the sound drivers that are loaded on the system.

Other Output Devices

Video cards and sound cards are two of the most popular types of output devices found in systems today, but they are not the only types of output devices. This section overviews some of the other types of output devices found in computer systems.

- ◆ **Multimedia cards:** Not really an output device, but when it comes to multimedia, video and sound files have to be stored somewhere. A number of devices today use *multimedia cards (MMC)*, which are sticks of memory the size of a stamp that are used with a multimedia device such as a music player or digital camera.
- ◆ **I/O devices:** Any number of input/output (I/O) devices may be connected to your system at any point in time. These devices, such as printers, scanners, and fax devices, can connect to your system in any form, such as SCSI, serial, or even parallel. Here is a summary of some of the technologies used with I/O devices.

- **SCSI:** SCSI is a technology that supports daisy-chaining — that is, having one device connect to the next. In order to connect a SCSI device to your system, you need a SCSI card, and you need to ensure that you assign unique IDs to each device and terminate both ends of the SCSI bus. Typically, storage devices such as drives, CD-ROMs, and tapes appear as SCSI devices.
- **Serial:** Some I/O devices are serial devices that connect to the serial port of the computer. An example of an I/O device that may connect to the serial port is a modem.

Remember that the serial connection is defined in the RS-232 standard and sends data one bit at a time.

- **Parallel:** Parallel devices get their name by receiving 8 bits of data at a time — as opposed to the 1 bit of data at a time that is supplied by serial connections. Parallel is defined in the IEEE 1284 standard and has different parallel modes, such as SPP, ECP, and EPP. SPP (standard parallel port) is unidirectional and transfers data at 150 Kbps. ECP (extended capabilities port) is bidirectional with a transfer rate of 2 Mbps. EPP (enhanced parallel port) is bidirectional and can transfer information at 2 Mbps — but also supports daisy-chaining!
- **USB:** USB is the most popular way to connect devices to your system today. USB 1.1 can transfer information at 12 Mbps, while USB 2.0 can transfer information at 480 Mbps. USB supports 127 devices in a single USB chain and is also hot-swappable — meaning that devices can be added and removed from the system without shutting the system down first.



Getting an A+

In this chapter, you find out about the two major types of output devices found on systems today — the video card and the sound card. The following summarizes some key points about video and sound cards:

- ◆ The video card converts digital data from the computer into analog data that is used by the monitor to create the display.
- ◆ VGA supports 16 colors at a resolution of 640 x 480.
- ◆ Super VGA supports 16 million colors at a resolution of 1280 x 1024.
- ◆ AGP cards are the popular form for video cards today, and sound cards are implemented as PCI cards.
- ◆ CRT monitors use an electron gun to stimulate the phosphor-covered screen and make it glow.

Prep Test

- 1** Sound cards are connected to the system using which of the following?
(Select two)
 - A AGP
 - B ISA
 - C PCI
 - D USB

- 2** You wish to enable multi-display support in Windows XP; how do you do this?
 - A Select the monitor and choose the Add Monitor option in the video card properties.
 - B Select the monitor and choose the Add Monitor option in the display properties.
 - C Select the monitor and choose the Extend My Windows Desktop to This Monitor option in the video card properties.
 - D Select the monitor and choose the Extend My Windows Desktop to This Monitor option in the display properties.

- 3** Which video standard supports 16 colors at 640 x 480 resolution?
 - A EGA
 - B VGA
 - C SVGA
 - D CGA

- 4** What IRQ is typically assigned to ISA sound cards?
 - A 1
 - B 4
 - C 5
 - D 7

- 5** What type of display is found on laptops?
 - A CRT
 - B LCD
 - C RCT
 - D USB

6 Which video standard supports only four colors?

- A EGA
- B VGA
- C SVGA
- D CGA

7 Which type of display uses an electron gun?

- A CRT
- B LCD
- C RCT
- D AGP

8 Which video standard supports 16 million colors at 640 x 480 resolution?

- A EGA
- B VGA
- C SVGA
- D CGA

Answers

- 1 B, C. Older sound cards were implemented as ISA cards, while newer sound cards are PCI cards. *See "Installing a sound card."*
- 2 D. To enable multi-display support, you go to the display properties in the Control Panel and choose the Extend My Windows Desktop to This Monitor option, located on the Settings tab. *Review "Configuring multi-display support."*
- 3 B. VGA supports 16 colors at 640 x 480 resolution. *Check out "Looking at the video standards."*
- 4 C. Sounds cards defaulted to IRQ 5 before Plug and Play became a standard. *Peruse "Types of sound cards."*
- 5 B. LCDs are found on laptops but are also used with desktop computers today. *Take a look at "LCD."*
- 6 D. CGA supports only four colors. *Peek at "Looking at the video standards."*
- 7 A. A CRT monitor uses an electron gun to light phosphors on the screen. *Look over "Cathode Ray Tube (CRT)."*
- 8 C. SVGA is the video standard used today, and it supports 16 million colors. *Study "Looking at the video standards."*

Chapter 4: Examining System Resources

Exam Objectives

- ✓ Identifying IRQs
- ✓ Identifying DMA channels
- ✓ Identifying I/O addresses
- ✓ Identifying available resources

A part of the A+ exam is all about troubleshooting and configuring a device during installation. In the past, a big part of troubleshooting computer systems was understanding how to configure resources such as IRQs, I/O addresses, and DMA channels on a device — which is what you find out to do in this chapter. However, due to Plug and Play operating systems, such as Windows XP, and Plug and Play devices, such as USB and PCI, you most likely won't find yourself configuring an IRQ or I/O address with today's systems. As a computer technician, though, you may support older hardware and older operating systems — CompTIA still has such topics listed in the A+ objectives, so we discuss them in this chapter to prepare you for the A+ hardware exam.

Suppose that one of your best friends, Bob, calls you up one day to borrow your perfectly functional sound card and you tell him to come over and pick it up. Bob comes over to your house to pick up the sound card, and you show him that the card works perfectly in your computer. You then remove it carefully from the system and hand it over to him. Bob takes the sound card home, inserts it into his system, and powers his system on. Windows boots up and installs the driver for the card, but after the driver has been loaded, Bob notices that there is no sound coming from the computer! Viewing the status of the sound card in the Windows Device Manager reveals that the device is not functioning properly. So Bob is thinking to himself, "Gee, the sound card worked perfectly on Glen's system. What gives?"

Or worse still, say you purchase a sound card from the local computer store, insert it into your computer, properly install the driver, and it doesn't work.

Problems like these are among the most common types of problems that A+ professionals will encounter when supporting older non-Plug and Play devices such as an old ISA card. In the past you had to be aware that a device that works in one computer doesn't initially work in another computer, even

when the operating systems on both computers are identical — the reason being that you needed to configure the device with the appropriate resources of that specific system.

This chapter helps you understand why many devices don't automatically work in your system and identify what you can do to make them work. You also discover the different types of resources that you will be required to troubleshoot and how to identify the default resources for any system.

Understanding System Resources

In this chapter, when you see the term *system resource*, you should think of it as a setting assigned to a device that allows it to function within the computer. A *device* is any piece of hardware that you can install on the computer: for example, a network card, a modem, or a sound card.

The three major system resources that can be assigned to devices are I/O addresses, IRQ addresses, and DMA addresses. A fourth system resource, called a memory address, can sometimes be assigned to devices as well. In the following sections, I discuss each of these system resources.

I/O addresses

It is extremely important to remember that the CPU is the traffic cop of the entire system. If something is going to happen on the system, the CPU (processor) usually enables the action. All devices in the computer need to communicate with the processor from time to time, and the processor needs a method of separating and prioritizing all these communications.

Because the processor needs to send information to a number of different devices, and because those devices need to know which messages coming from the CPU are for them, each device is assigned an *I/O address*, or *input/output address*. The I/O address is a special port address that represents a pathway between the CPU and the device. So, for example, if the processor needs to send information to LPT1, it can send the information to pathway 0378–037F, which is the pathway address that leads to LPT1. (For more on the standard I/O address assignments, see Table 4-1, later in this chapter.) Think of these pathways as tunnels; each device has its own tunnel that extends from the device to the processor and is used for communication between the device and the processor.

Figure 4-1 shows a number of devices that are configured with I/O addresses. (Note that the addresses in the figure are not necessarily the actual addresses that are used by those devices — I made up these addresses simply to get the point across.) In this example, if the processor needs to send information to the sound card, it knows that if it sends the information down I/O port address

220, then the sound card will receive the information. Conversely, when the processor receives information from I/O port address 220, it knows that the information came from the sound card because the address of 220 is assigned to only one device — the sound card!

There are 65,536 different I/O port addresses available on the system. (There are actually fewer *usable* addresses than that because when you assign an I/O address to a device, you are really assigning a range of addresses.) The trick to troubleshooting resources is to make sure that you have not assigned the same I/O port address to two different devices. If you do, you will get a *resource conflict*. A resource conflict is when two devices are using the same resource, such as an I/O address, IRQ, or DMA channel.



The I/O address assigned to a device is really a range of values, such as 0378 through 037F. Most people refer to the address block by the first value in the range — in this example, 0378.

To prevent resource conflicts, each device should have a unique I/O address. But how do you know which I/O addresses are already being used by existing devices? One way is to use Device Manager to view I/O addresses being used on the system.




Figure 4-1: Each device uses a unique I/O address to send and receive information to and from the processor.

When troubleshooting Windows 2000, Windows XP, or Windows 2003, you can use Device Manager to obtain a list of resources in use. To view I/O addresses currently in use with these operating systems, follow the steps below:

1. Right-click My Computer and choose Properties.
2. Click the Hardware page tab.
3. Click the Device Manager button.
4. Choose View → Resources by Type.

The screen should look similar to the one shown in Figure 4-2.

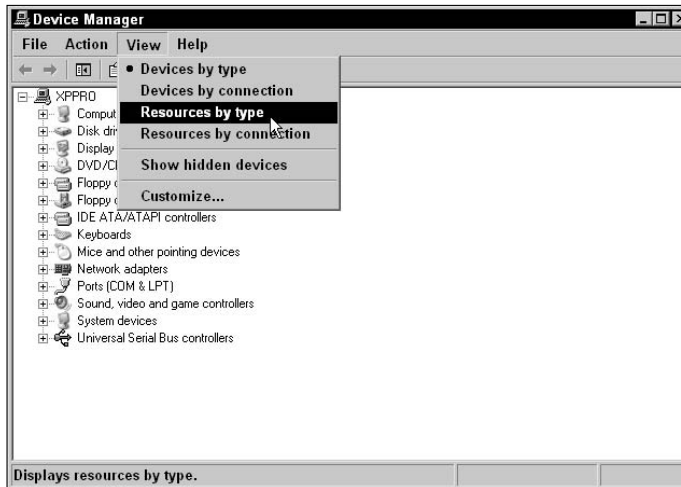


Figure 4-2:
Configuring
Device
Manager
to view
resources in
Windows
XP.

5. Expand Input/Output (I/O) and view the list of I/O addresses being used (shown in Figure 4-3).

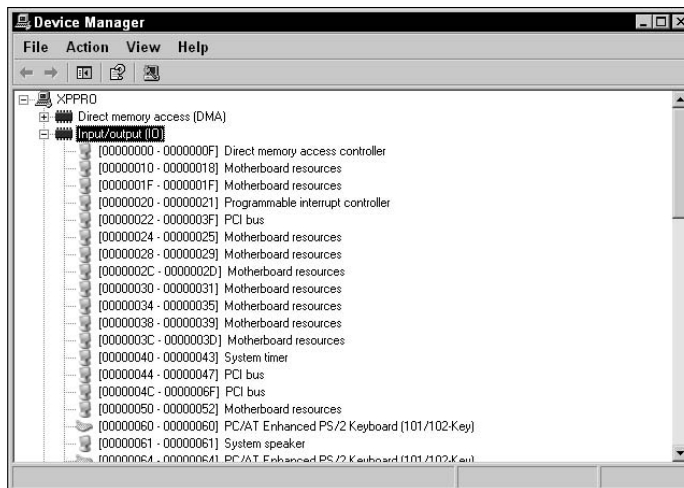


For the A+ hardware exam, it is important to memorize the I/O addresses of standard ports, such as COM1, COM2, and LPT1. Table 4-1 lists the I/O addresses you are required to know for the A+ hardware exam.

Device	I/O Address Range
COM1	03F8 to 03FF
COM2	02F8 to 02FF
COM3	03E8 to 03EF
COM4	02E8 to 02EF

<i>Device</i>	<i>I/O Address Range</i>
LPT1	0378 to 037F
LPT2	0278 to 027F
Math coprocessor	00F8 to 00FF
Primary hard disk controller	01F0 to 01F7
Secondary hard disk controller	0170 to 0177
Sound cards	0220 to 022F
Floppy disk	03F0 to 03F7

Figure 4-3:
Viewing I/O
Addresses
in use with
Windows
XP Device
Manager.



Interrupt ReQuest (IRQ)

Each device has its own tunnel for sending and receiving information to and from the processor, which is the function of the I/O port address discussed in the previous section. But the processor is busy doing something important nearly all the time. How does the processor know when a device has information to communicate to the processor? Too much overhead would be created if the processor had to continuously poll each device to see whether it had data to send to the processor; instead, each device is responsible for notifying the processor if it has information to send. Because devices are responsible for notifying the processor that they have information, the devices need a way to interrupt the processor from its current work to ask it to service their requests. The method that is used to interrupt the processor is called an *Interrupt ReQuest*, or *IRQ* line.

If you were standing beside someone who was involved in a conversation and you really wanted to talk to that person, what would you do? Most people would tap the person on the shoulder and say something like, “Excuse me, sir, can I send you data?” Okay, maybe you wouldn’t say that, but you get the idea. Tapping the person on the shoulder is similar to what a device does with the processor when it wants to send the processor some information. The device sends a signal down the IRQ line, which virtually taps the processor on the shoulder — as a result, the device grabs the processor’s attention and the processor is ready to receive information via the I/O address.

When a device taps the processor on the shoulder, the processor needs to know which device needs attention. This is why each device is assigned a unique IRQ line number. When a device sends a signal down the IRQ line to interrupt the processor, the processor checks which line the signal originated from and then attends to that device. For example, in Figure 4-4 (remember that the values shown in the figure are for discussion purposes only — use Tables 4-2 and 4-3, later in this section, to find out which IRQs are used by devices to prepare for the exam), if the network card wants to send information to the processor, the network card must first get the processor’s attention by sending a signal down IRQ 10.




Figure 4-4: IRQs are how a device grabs the CPU’s attention.



It is important to note that when information is sent to the processor, it is sent through the I/O address (the tunnel), not the IRQ line. So the IRQ just grabs the processor's attention while the I/O address is used for the actual delivery of the information.

Cascading IRQs

Originally there were only 8 IRQs available on XT (before 286) systems, but now there are 16 IRQs available on AT (after 286) systems. In order to get 16 IRQs, another IRQ controller was added to the system. Having two sets of IRQs managed by two different controllers presented some technical problems, so to help the two IRQ controllers act as one unit, they are cascaded (or linked) together by IRQ 2 and IRQ 9. The second controller goes through the first controller to send requests. Figure 4-5 shows the two controllers linked together.

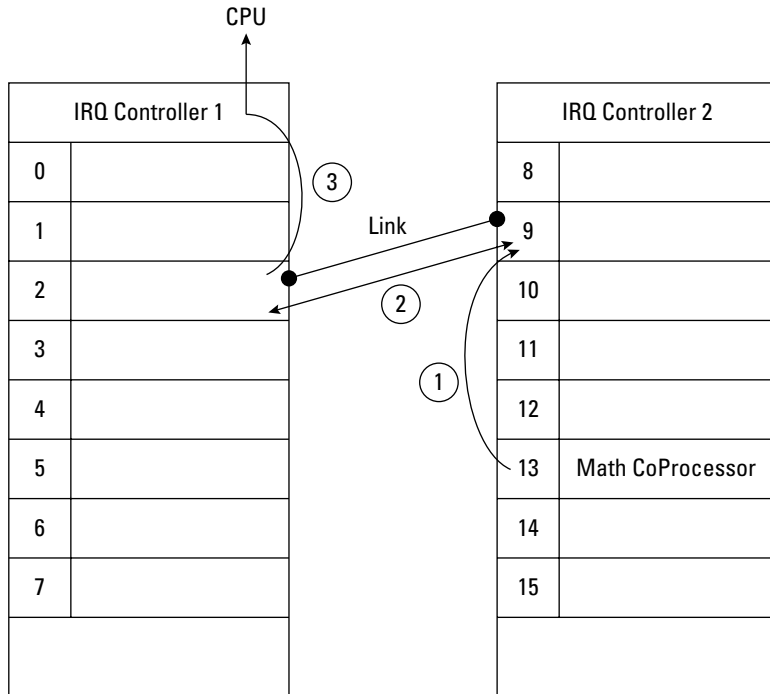


Figure 4-5: Two IRQ controllers are cascaded together.

Here's how cascading controllers work. In Figure 4-5, the path that the math coprocessor (which uses IRQ 10) takes to send an interrupt request to the processor is shown in three steps:

1. The math coprocessor attempts to send a signal to the CPU, but because the second interrupt controller handles the math coprocessor, the signal is passed to IRQ 9, which forwards all signals for the second interrupt controller.
2. IRQ 9 passes the signal to its linked partner, IRQ 2, which is managed by the first interrupt controller.
3. IRQ 2 then passes the signal through the first interrupt controller, which then sends the signal to the CPU.



Because older systems only had one set of eight IRQs, when the additional eight IRQs came along only the first interrupt controller was allowed to send information to the CPU. This means that when interrupts 8 through 15 send an interrupt request, they must pass it to the first controller, which in turn passes it to the processor on their behalf.

Like I/O addresses, if you assign two devices the same IRQ value, you get a resource conflict that results in at least one (and maybe both) of the devices not working. To prevent assigning two devices the same resource, you need to understand which IRQ values are already being used by your system.

To view a list of IRQs in use on a Windows 2000, Windows XP, or Windows Server 2003 system, follow these steps:

1. **Right-click My Computer and choose Properties.**
2. **Click the Hardware page tab.**
3. **Click the Device Manager button.**
4. **Choose View⇨Resources by Type.**
5. **Expand Interrupt Request (IRQ) and view the list of IRQs being used.**

The screen should look similar to the one shown in Figure 4-6.



I should note that due to the low number of IRQs that are available on a system, Intel has been incorporating a newer technology in its new processors known as *Advanced Programmable Interrupt Controller (APIC)*. Intel currently has APIC supporting up to 24 IRQs.



For the exam, it is important to memorize the standard IRQs that are in use and which IRQs are generally considered available. Table 4-2 shows a listing of the standard IRQs and their use. Be sure to memorize this table for the A+ hardware exam.

Figure 4-6:
Viewing
IRQs in
use with
Windows
XP Device
Manager.

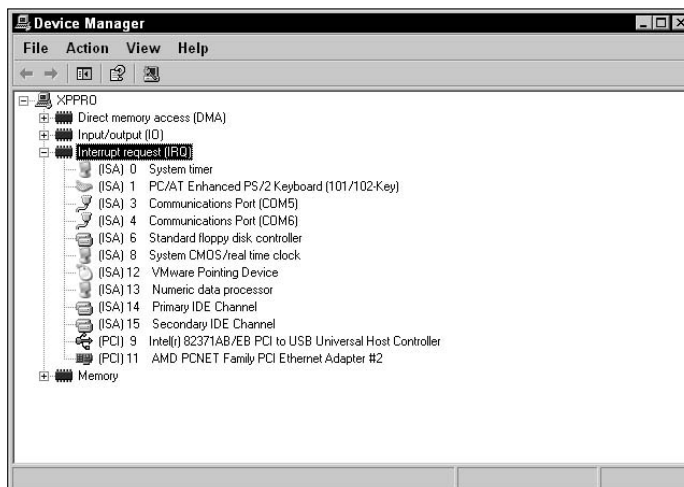


Table 4-2 Standard IRQ Assignments

<i>IRQ Value</i>	<i>Device</i>
0	System Timer
1	Keyboard
2	Link to second IRQ controller
3	COM2, COM4
4	COM1, COM3
5	LPT2
6	Floppy disk drive
7	LPT1
8	Real time clock
9	Available, but should not be used if IRQ 2 is being used
10	Available
11	Available
12	Available if not used by PS/2 mouse
13	Math Coprocessor
14	Hard disk controller
15	Available

For the exam, it is important to know that the lower the IRQ value, the higher the priority the device will have with the processor. Also, be sure to memorize the default assignments of IRQs and I/O addresses for the exam.



Here are a few important points about IRQ assignments:

- ◆ **IRQs 10, 11, 12, and 15 are generally available.** If you are installing a new device into a computer and need to assign an IRQ, try one of these IRQ values first.
- ◆ **IRQ 3 and IRQ 5 are used by COM2 and LPT2, respectively.** If you are not actually using COM2 or LPT2, you can consider IRQ 3 and IRQ 5 as being available.



Table 4-3 contains a listing of ports and associated addresses to memorize for the exam. I recommend memorizing this table and then, when you sit down to take the A+ hardware exam, immediately writing this down on scrap paper before you start the exam. You are not allowed to take paper with the table listing into the exam with you, but nothing stops you from making the table on your scrap paper before you start the exam and then referring to it with each question on this topic area.

Table 4-3		Default IRQs and I/O Addresses
<i>Device</i>	<i>IRQ</i>	<i>I/O Address</i>
COM1	4	3F8-3FF
COM2	3	2F8-2FF
COM3	4	3E8-3EF
COM4	3	2E8-2EF
LPT1	7	378-37F
LPT2	5	278-27F



You may be asked on the exam the IRQ of a modem plugged into a COM port. Know that an IRQ is not really assigned to a modem, but to the COM port the modem is using. You will also want to have an idea of where infrared and USB devices will sit as far as IRQs are concerned. Infrared ports attach themselves to virtual COM or LPT ports, which means that your infrared device will probably end up using IRQ 3 or 4. USB ports use the IRQ associated with the PCI bus, which can be anywhere from IRQ 9 on up.

Direct Memory Access (DMA)

A number of different devices today require constant access to system memory. Normally, devices must go through the CPU to write information to system memory, but using such a scheme can cause a lot of unnecessary overhead, so why not allow a device to access memory directly?

To increase performance and to offload some of the work from the CPU, you can assign some devices a *Direct Memory Access (DMA)* channel. The DMA channel is a special pathway that allows the device to read and write information directly to system memory without passing the data to the processor.

For example, in Figure 4-7 you can see that the modem has been assigned DMA channel 6 and the printer has been assigned DMA channel 5. (Like I/O addresses, these addresses may be different on each system. Remember that the values I use in the figure are fictitious — again, refer to Tables 4-1, 4-2, and 4-3 when preparing for the exam.) In this example, the modem and the printer can write information to memory directly, whereas the network card and the sound card must pass through the CPU.

There are only eight DMA channels available on your system, which should not be a huge problem because not all devices use DMA channels. Some examples of the different devices that you may run into that use DMA channels are sound cards, network cards, and occasionally CD-ROM drives. Table 4-4 shows a listing of common DMA channels.




Figure 4-7: DMA channels are used to give a device direct access to memory.

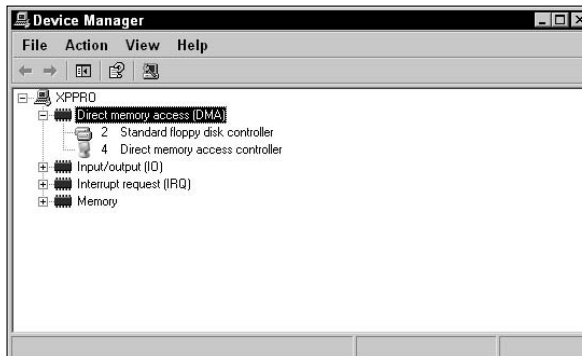
DMA Channel	Device
0	Available
1	Sound or Available
2	Floppy Drive
3	Available
4	Cascade
5	Sound or Available
6	Available
7	Available

Like IRQs, two DMA controllers are linked by a cascading DMA channel, DMA channel 4. DMA channels 0–3 are for 8-bit boards and cards; DMA channels 5–7 are used for 16- and 32-bit cards.

To view the DMA channels that are in use on your system, you can use the Windows *Device Manager* utility. The following steps demonstrate how to view the DMA channels in use within Windows 2000/XP and Windows Server 2003 systems:

- 1. Right-click My Computer and choose Properties.**
- 2. Click the Hardware tab.**
- 3. Click the Device Manager button.**
- 4. Choose View ⇨ Resources by Type.**
- 5. Expand Direct Memory Access (DMA) and view the list of DMA channels being used (shown in Figure 4-8).**

Figure 4-8:
Viewing
DMA
addresses
in use with
Windows
XP Device
Manager.



Memory addresses

A less common resource that may be assigned to devices is a *memory address*. A memory address is an area of memory where the device is allowed to store information.

If multiple devices have been assigned access to the same memory address, a device conflict will occur, and one or both devices may not function. When troubleshooting devices, you need to look out for memory address conflicts. The following steps demonstrate how to view memory addresses in use within Windows 2000/XP and Windows Server 2003.

1. Right-click My Computer and choose Properties.
2. Click the Hardware page tab.
3. Click the Device Manager button.
4. Choose View → Resources by Type.
5. Expand Memory and view the list of memory addresses being used (shown in Figure 4-9).

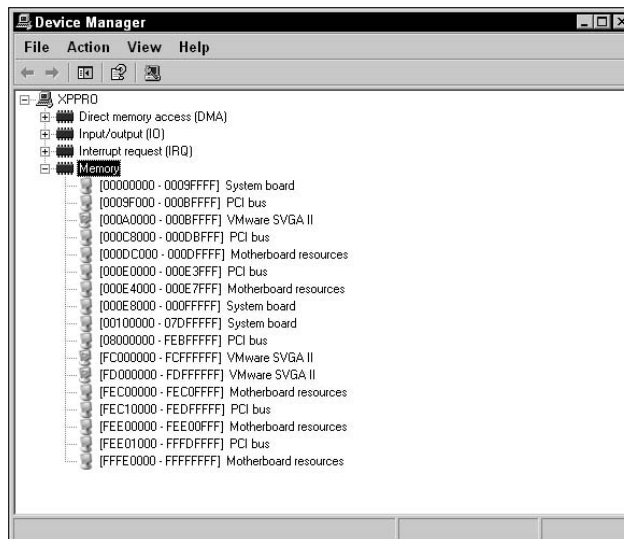


Figure 4-9:
Viewing
memory
addresses
in use with
Windows
XP Device
Manager.

Working with System Resources

In the preceding sections, you find out about system resources and how to view them on a system. In the following sections, I introduce you to methods of viewing and editing system resource settings. One of the biggest issues in

troubleshooting device-installation problems is being able to identify and solve resource conflicts.

Identifying resource conflicts

You can use a number of different tools to view system resources (like I/O addresses, IRQs, and DMA channels) and determine whether a resource conflict exists. Some of the tools are shipped with your operating system, and other tools are third-party products that you need to purchase.

Device Manager is a Windows utility that you can use to view system resources and identify problems with devices. One of the benefits of the Windows Device Manager is that it not only shows you a device that has a conflict, but also the device that it is in conflict with. To view conflicts with a resource on the device, right click on the device in Device Manager and choose Properties. Once in the properties of the device, click on the resource page tab. When you select a resource, you will see any conflicts at the bottom of the window (shown in Figure 4-10).

Notice in Figure 4-10 that the LPT port is the device you are troubleshooting. This device has an I/O conflict because you have assigned the LPT port the same I/O address as the ISAPNP Read Data Port. Notice also that the conflicting I/O address is 0278–027F.

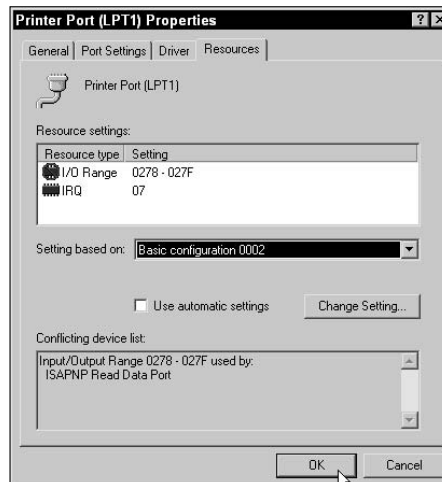


Figure 4-10:
Identifying a
conflicting
device with
Device
Manager.

With today's systems, this conflict wouldn't happen because most devices are PCI or USB devices and fully support Plug and Play. This means that

when the operating system starts up, it dynamically assigns each device a unique resource. However, if you are supporting older devices and operating systems, you may come across conflicts.

Changing system resources

Although viewing system resources and identifying resource conflicts are important, you also need to be able to make changes that will solve device conflicts.

This section discusses troubleshooting techniques and different ways of changing system resources. You find out different ways to configure resources such as I/O addresses, IRQs, and DMA channels. You also discover the three typical ways to configure a device for a resource: through jumper settings, setup disk, or Plug and Play.

Troubleshooting techniques

Before changing a system resource to solve a device conflict, you need to have a strategy for deciding the new value to assign to the resource. In general, the strategy is to look at all available resources and then reassign them such that each device uses a unique resource address.



With today's Plug and Play systems, you normally don't need to configure I/O addresses or IRQs because the system automatically configures those resources on the device when the device is installed.

Following is a scenario that should help you understand how to troubleshoot resource conflicts. Before Plug and Play became popular, the manufacturer of a device would limit the system resources that could be assigned to the device by providing a setup program that you would run and limit the IRQs that appeared in the setup program. In Figure 4-11, the sound card and network card work with only three IRQs each. The network card gives you the choice of using IRQ 4, 5, or 10; the sound card can function with IRQ 3, 4, or 5.

Network Card	Sound Card
IRQ 4*	IRQ 3*
IRQ 10	IRQ 4
IRQ 5	IRQ 5

Figure 4-11:
An example
of trouble-
shooting
IRQs, Part 1.


* = Assigned IRQ For That Device

Suppose that you assign IRQ 4 to the network card and IRQ 3 to the sound card. These resources are normally used by COM1 and COM2, but because you are not using COM1 and COM2, you decide to use those IRQs — makes sense!

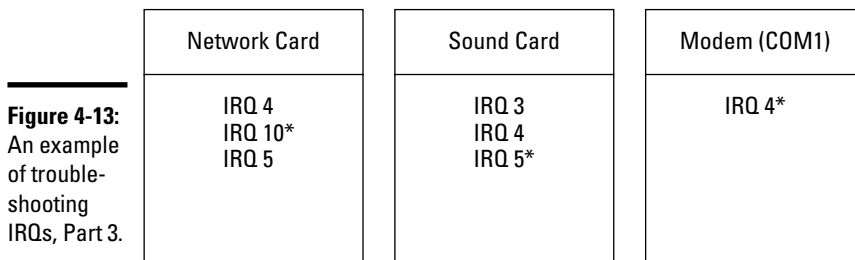
A few months later, you decide that you want to install an external modem on the computer. You connect the modem to COM1, which uses IRQ 4, forgetting that you have already assigned IRQ 4 to the network card. Connecting the modem to COM1 generates a resource conflict that causes the network card and the modem not to work! Now the fun begins . . .

To fix such a resource conflict, you need to look at the IRQs that can be used for each device and then start juggling resource assignment around until you have all three devices working together on the same system. Figure 4-12 shows a modem being added to the scenario.

You can see in Figure 4-12 that when the modem is installed, it will use IRQ 4 and conflict with the network card that is also using IRQ 4. At the same time, you notice that your sound card uses IRQ 3 (normally used by COM2), so to prevent future conflicts, you may be better off reassigning the IRQ for the sound card as well. You decide to assign the network card to IRQ 10 so that the sound card can be assigned IRQ 5. The outcome of these reassignments is shown in Figure 4-13.



* = Assigned IRQ For That Device



* = Assigned IRQ For That Device

The final outcome of your decision on IRQ assignments allows all three devices to function without conflicts, and you are prepared for the day that you decide to use COM2 by having IRQ 3 free.

Changing non-Plug and Play resources

If you are supporting non-Plug and Play devices, such as old ISA devices, you need to manually configure each non-Plug and Play device with system resources manually. Changing system resources on such a device may take a little more research than managing these resources in a Plug and Play environment. There are two popular methods of changing resources in non-Plug and Play devices: Modify the setting through the use of a jumper, or use the software setup program that came with the device.

A *jumper* (shown in Figure 4-14) is a set of pins on a card or board that applies a certain setting when a circuit is closed on a certain subset of the pins. Each jumper set on a card has a label associated with it that looks something like *J10*, for *Jumperset 10*. Exactly what this label looks like is, of course, up to the manufacturer of the device.




Figure 4-14:
Jumpers
are used to
configure
resources.

The documentation (finding the documentation is the challenge!) for the card tells you what the jumper set of J10 is actually used for. Suppose that the jumper set of J10 sets the IRQ value, and that the jumper set of J10 has

6 pins. The documentation for the card may tell you, for example, that you need to place the jumper over pins 1 and 2 to assign IRQ 4, across pins 3 and 4 to assign IRQ 5, and across pins 5 and 6 to assign IRQ 10.



You should keep all documentation for your system and any devices you purchase so that any information about the device is available when needed. If you misplace the documentation, you can usually find the information on the manufacturer's Web site.

Another method to change the values of system resources is through *software setup*, which allows you to change the I/O address, IRQ, or other settings by running a setup program. The setup program usually comes on a floppy disk or CD-ROM and is a much simpler process than messing around with jumper settings (now you just have to worry about not losing the setup disk!). Again, if you misplace the setup disk, you can normally get the program from the manufacturer's Web site.



It is a good idea to make a backup copy of the setup disk that comes with each device in case something happens to the original disk.

When you run the setup program for the device, you usually have options to view and change the system resources. You change the system resources by saving the setup information (also known as *flashing*) to the EEPROM chip on the card (usually by pressing the F10 key while in the setup program). Figure 4-15 shows an example of a setup (diagnostic) program used to modify the configuration of an old ISA network card.

Figure 4-15:
Changing resources on a non-Plug and Play device using the setup disk.



Changing Plug and Play resources

In a Plug and Play environment, changing settings is a bit easier than in a non-Plug and Play environment. *Plug and Play* systems automatically assign resources to devices dynamically as the operating system starts up. In order to

have a full Plug and Play environment, you must have a Plug and Play operating system, such as Windows, a Plug and Play device, and a Plug and Play BIOS.

In Windows XP, if the operating system has wrongly assigned a resource to a Plug and Play device, you can change it through Device Manager:

1. **Right-click My Computer and select Properties.**
2. **Click the Hardware tab.**
3. **Click the Device Manager button.**

The Device Manager utility appears.

4. **In the Device Manager utility, right click on the device and choose properties.**

Doing so shows you the properties of the device.

5. **Click the resources page tab. If you are allowed to change the resources on the device, uncheck the Use Automatic Settings check box.**

After you uncheck the Use Automatic Settings check box, you may choose a different resource configuration.

6. **Select a new configuration setting from the Settings Based On dropdown list (shown in Figure 4-16).**

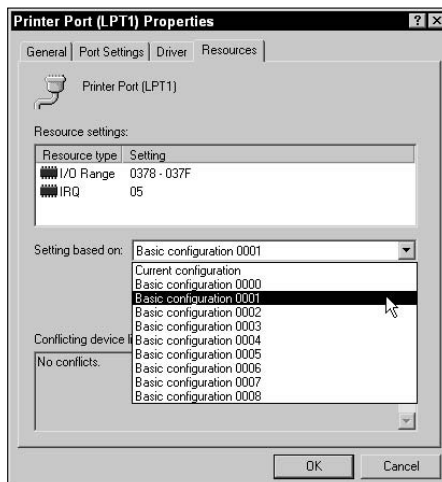


Figure 4-16:
Changing
resources
through
Device
Manager in
Windows
XP.

After choosing a different setting you will notice the resource change above in the resource settings portion of the dialog box.

7. **Choose OK and then close Device Manager.**

Reserving resources

You may encounter an older system that has a mixture of Plug and Play devices and legacy non-Plug and Play devices, which can present problems because the legacy devices have been configured with specific resources, and you want to ensure that the operating system doesn't assign those resources to Plug and Play devices. To prevent the operating system from assigning the same resource to a Plug and Play device, you can *reserve* that resource.

Reserving the resource can be done at two different levels, as the following list details:

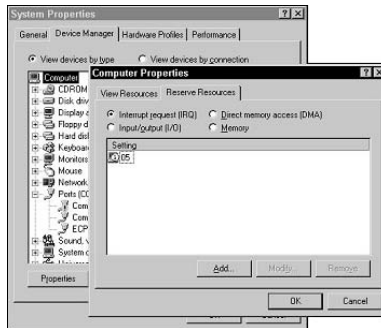
- ◆ **Reserving a resource at the hardware level:** Reserving a resource at the hardware level means that you enter the CMOS setup program and tell the system what resources it is not allowed to give out (see “Reserve Resources in Book II, Chapter 4).
- ◆ **Reserving a resource at the software level:** You can also reserve resources at the software level by using Device Manager in older versions of Windows. If you configure the resource at the software level, the resource is reserved only for the operating system you put the reservation in for. So installing a new operating system means you need to reserve the resource again.

The following steps illustrate how to reserve resources in Windows 9x operating systems:

1. Choose **Start**⇨**Settings**⇨**Control Panel**.
2. Double-click the **System** icon to display the **System Properties** dialog box.
3. Click the **Device Manager** tab.
4. Click **Computer** at the top of the device list and then click the **Properties** button.
5. Click the **Reserve Resources** tab (as shown in Figure 4-17).
6. Select the type of resource that you want to reserve (**IRQ**, **I/O**, **DMA**, or **Memory**) and then click the **Add** button.
7. In the **Edit Resource Setting** dialog box, type the address you want to reserve for the resource and then click **OK**.

When working with new operating systems, such as Windows XP, and newer devices, such as PCI and USB, the operating system does not allow you to reserve a resource because you would not really need to. You should allow the operating system to manage the environment through Plug and Play. If you find a reason to reserve the resource when using a newer operating system, you can resort to reserving the resource in CMOS.

Figure 4-17:
Reserving
resources
with Device
Manager.



To practice working with system resources, take a look at Lab 4-1, Lab 4-2, Lab 4-3, and Lab 4-4. These labs can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Getting an A+

This chapter introduces one of the most heavily tested topics with servicing computers: Managing resources. The following is a list of key points to remember when troubleshooting system resources:

- ◆ An *I/O address (input/output address)* is the pathway used by a device to send and receive information to and from the processor.
- ◆ An *IRQ (Interrupt ReQuest)* is a signal that a device can send to interrupt the processor from its current task.
- ◆ A *DMA (Direct Memory Access) channel* is a device's direct pathway to memory so that it can read and write information to memory without bothering the processor.
- ◆ Changing resources on *legacy devices* usually involves changing jumper settings or running the diagnostic program for the device.
- ◆ Changing resources on *Plug and Play devices* involves going to Device Manager in Plug and Play operating systems and changing the resource for the device.

Prep Test

- 1** Device A has been configured for IRQ 3. Which of the following devices will device A potentially conflict with?
 - A COM1
 - B LPT1
 - C COM2
 - D LPT2

- 2** What is the default I/O address assigned to COM1?
 - A 02F8
 - B 0378
 - C 03F8
 - D 0278

- 3** Which of the following best describes DMA?
 - A DMA stands for Dynamic Memory Address, which means that the memory address is dynamically assigned to the device each time the system reboots.
 - B DMA stands for Direct Memory Access, which means that a device can use the operating system's area of protected memory without seeking permission.
 - C DMA stands for Dynamic Memory Address, which means that the device is assigned resources dynamically as the system boots.
 - D DMA stands for Direct Memory Access, which means that the resource is given a special pathway to memory so that it does not have to go through the processor.

- 4** You have assigned your network card to I/O address 0378. Which device will the network card create a conflict with?
 - A COM1
 - B COM2
 - C LPT1
 - D LPT2

- 5** What IRQ is assigned to COM1?
 - A 3
 - B 5
 - C 7
 - D 4

- 6** What is the default I/O address assigned to COM2?
- A 02F8
 - B 0378
 - C 03F8
 - D 0278
- 7** In order for a device to interrupt the CPU and request service, the device must be assigned what?
- A IRQ
 - B I/O address
 - C DMA channel
 - D Memory address
- 8** What is the default I/O address assigned to LPT1?
- A 02F8
 - B 0378
 - C 03F8
 - D 0278
- 9** What IRQ is assigned to COM2?
- A 3
 - B 4
 - C 7
 - D 5
- 10** Which two ports use IRQ 4?
- A COM1
 - B COM2
 - C COM3
 - D COM4
- 11** What is the default I/O address assigned to LPT2?
- A 02F8
 - B 0378
 - C 03F8
 - D 0278

Examining System Resources

- 12** What IRQ is assigned to LPT1?
- A 2
 - B 4
 - C 5
 - D 7
- 13** What two ports use IRQ 3?
- A COM1
 - B COM2
 - C COM3
 - D COM4
- 14** What IRQ is assigned to LPT2?
- A 2
 - B 4
 - C 5
 - D 7
- 15** How many I/O port addresses are there in total?
- A 8
 - B 16
 - C 16,384
 - D 65,536
- 16** What IRQ value is used to cascade the first 8 IRQs to the second 8 IRQs?
- A 0
 - B 1
 - C 2
 - D 7
- 17** How many DMA channels are there in total?
- A 8
 - B 16
 - C 16,384
 - D 65,536
- 18** How many IRQs are there on an AT system?
- A 2
 - B 5
 - C 8
 - D 16

19 Which Windows utility can you use to view system resources in use?

- A Windows Explorer
- B Device Manager
- C Internet Explorer
- D Regedit

20 How many IRQs are there on an XT system?

- A 5
- B 8
- C 16
- D 32

Answers

- 1 C. If device A is assigned IRQ 3, it will conflict with anything connected to COM2 because COM2 (as well as COM4) uses IRQ 3 by default. COM1 and COM3 use IRQ 4 by default, LPT1 uses IRQ 7, and LPT2 uses IRQ 5 by default. *See "Interrupt ReQuest (IRQ)."*
- 2 C. The default I/O address of COM1 is 03F8. The default I/O address of COM2 is 02F8, the default I/O address of LPT1 is 0378, and the default I/O address of LPT2 is 0278. *Review "I/O addresses."*
- 3 D. DMA (Direct Memory Access) is a device's unique path to memory so that it can read and write to memory quickly. The benefit of a device receiving a DMA channel is that it will not have to pass through the CPU each time it needs to access a block of memory. *Check out "Direct Memory Access (DMA)."*
- 4 C. LPT1 uses the default I/O address of 0378, so installing a network card and configuring it for the same address will create an I/O address conflict. The default I/O address of COM1 is 03F8. The default I/O address of COM2 is 02F8, and the default I/O address of LPT2 is 0278. *Peruse "I/O addresses."*
- 5 D. COM1 defaults to IRQ 4. IRQ 3 is used by COM2, IRQ 5 is used by LPT2, and IRQ 7 is used by LPT1. *Take a look at "Interrupt ReQuest (IRQ)."*
- 6 A. COM2 uses the default I/O address of 02F8. LPT1 uses the default I/O address of 0378, COM1 uses the I/O address of 03F8, and LPT2 uses 0278. *Peek at "I/O addresses."*
- 7 A. IRQs interrupt the CPU and request service from the CPU. This is where the acronym IRQ (Interrupt ReQuest) comes from. An I/O address is a device's communication channel to the CPU so that it can send and receive information to and from the CPU. A DMA channel is a device's direct path to memory so that it can quickly read and write to memory. A memory address is an area of memory used by the device to store information. *Look over "Interrupt ReQuest (IRQ)."*
- 8 B. The default I/O address assigned to LPT1 is 0378. The default I/O address assigned to LPT2 is 0278, the default I/O address of COM1 is 03F8, and the default I/O address of COM2 is 02F8. *Study "I/O addresses."*
- 9 A. IRQ 3 is the default IRQ of COM2. IRQ 4 is the default IRQ of COM1, while IRQ 7 is the default IRQ of LPT1. IRQ 5 is the default IRQ of LPT2. *Refer to "Interrupt ReQuest (IRQ)."*
- 10 A and C. The two ports that use IRQ 4 are COM1 and COM3. COM2 and COM4 use IRQ 3. *Examine "Interrupt ReQuest (IRQ)."*
- 11 D. The default I/O address of LPT2 is 0278. LPT1 uses 0378, COM1 uses 03F8, and COM2 uses 02F8 as the default I/O address. *See "I/O addresses."*

- 12** D. LPT1 has the default IRQ of 7. IRQ 2 is the cascade IRQ that creates a link to the second IRQ controller. IRQ 4 is used by COM1 and COM3, while IRQ 5 is used by LPT2. *Review "Interrupt ReQuest (IRQ)."*
- 13** B and D. The two ports that use IRQ 3 are COM2 and COM4. The ports of COM1 and COM3 use IRQ 4. *Check out "Interrupt ReQuest (IRQ)."*
- 14** C. The default IRQ of LPT2 is IRQ 5. IRQ 2 is used to create a link to IRQ 9, which is cascaded to the second IRQ controller. IRQ 4 is used by COM1 and COM3, and IRQ 7 is used by LPT1. *Peruse "Interrupt ReQuest (IRQ)."*
- 15** D. There are 65,536 I/O address ports available on the system. There are 8 DMA channels available and 16 IRQs available. *Take a look at "I/O addresses."*
- 16** C. The IRQ that creates a cascade link to the second set of IRQs managed by the second IRQ controller is IRQ 2. The system timer uses IRQ 0, while IRQ 1 is used by the keyboard. IRQ 7 is used by LPT1. *Peek at "Interrupt ReQuest (IRQ)."*
- 17** A. There are 8 DMA channels on systems today. There are 16 IRQs and 65,536 I/O addresses available. *Look over "Direct Memory Access (DMA)."*
- 18** D. An AT system uses 16 IRQ channels that are numbered from 0 to 15. There were originally only 8 IRQ channels on XT systems, but a second set of 8 was created for AT systems. *Study "Interrupt ReQuest (IRQ)."*
- 19** B. Windows Device Manager is the tool that allows you to view system resources like IRQs and I/O addresses. You can also use Device Manager to help solve resource conflicts. *Refer to "Identifying resource conflicts."*
- 20** B. There are only 8 IRQs available on XT systems, but an additional set of 8 was created for AT (after 286 processors) systems, creating a total of 16 IRQs for AT systems. *Examine "Interrupt ReQuest (IRQ)."*

Chapter 5: Managing Printers

Exam Objectives

- ✓ Identifying paper feeder mechanisms
- ✓ Understanding types of printers
- ✓ Identifying printer connections
- ✓ Identifying common printer problems
- ✓ Understanding preventative maintenance and safety precautions

In today's business world, maintaining a printing environment is one of the most time-consuming parts of managing a network. For a supposedly paperless era, we seem to spend a lot of time troubleshooting why the printer won't print!



This chapter introduces you to the different types of printers and describes how each type of printer works, which is important for the A+ Certification exam. From an exam point of view, be sure you are extremely comfortable with the parts of a laser printer and the six phases of the laser printer's print process — you are guaranteed to get some questions on it!

Be sure to also go over the troubleshooting sections. I think that most people simply skim over this topic area, but you will be presented with some common problems on the exam, and you are responsible for identifying which component of the printer is causing the problem. Be sure to spend some time on this chapter!

Learning about Paper Feeder Mechanisms

A printer is useless without something to print on, so how the paper gets into a printer seems like a perfect place to start the discussion. The paper enters the printer through a *feeder*, technically a *paper feeder mechanism*. The paper feeder mechanism pulls or pushes the paper into the printer. There are two types of paper feeder mechanisms:

- ◆ **Continuous form feed**
- ◆ **Friction-feed**

The following sections take a close look at both of these types of paper feeders.

Continuous form feeders

Continuous form feed printers use a continuous sheet of 8.5-inch- or 14-inch-wide paper. The continuous sheet feeds through the printer, and then individual sheets are separated after the print operation — the pages are usually perforated to make this easier and to maintain a consistent size.

The paper has holes along both sides that fit over the sprockets located on the feeder wheel. When the wheel turns, it feeds the paper into the printer. The holes are torn off the paper after the printout is complete.

You may be wondering what types of printers still use this type of feeder mechanism. One common use of continuous feed printers is for printing company paychecks. Blank checks are connected as a continuous sheet of paper and fed into the printer. After a check prints, it moves out of the printer and the next check is fed in. After all the checks have printed, the accountant then separates each check for distribution to the employees.

Dot matrix printers are an example of a continuous form feed printer.

Friction feeders

The most popular type of printers are friction feed printers. A friction feed printer uses two rollers that pick up individual sheets of paper from a paper tray and feed them through the printer.

With friction feed printers, many different types of rollers are used throughout the print process. The rollers that pick the paper out of the tray are called “pickup” rollers, but other rollers pass the paper through the different parts of the printer.

Laser printers, photocopiers, and fax machines are friction feed printers.

Understanding Types of Printers



You'll find three different types of printers in today's busy world of computing: laser, inkjet, and dot matrix. You need to be familiar with each type for the A+ exam.

Laser printers

The *laser printer*, also known as a *page printer* because it prints one page at a time, is the most popular type of printer because it is fast, reliable, and has the best-quality printout of the three types of printers.

The laser printer gets its name because it uses a laser beam in the printing process. The laser printer, shown in Figure 5-1, is also the most expensive type of printer due to its high-cost components such as the laser.

Many parts work together to make the laser printer and its printing process run smoothly:

- ◆ **Paper feeder mechanism:** Laser printers use a set of pickup rollers to grab the paper from the paper tray and feed it into the printer.
- ◆ **Paper transport path:** Rollers are used throughout the print process so that the paper can continue to move through the printer.
 - *Registration rollers* move the paper.
 - *Fuser rollers* melt the toner onto the paper.
 - *Exit rollers* guide the paper out of the printer.



Figure 5-1:
Looking
at a laser
printer.

- ◆ **Toner cartridge:** The toner cartridge (shown in Figure 5-2) contains the replaceable components of the printing process. It contains three core components:
 - *The toner* is electrically charged material made up of pigment (to give it its color) and plastic (so it can be melted to the page) that is attracted to the paper to create the printout.
 - *The print drum* holds an electromagnetic charge when exposed to the laser. That charge then attracts the toner to the page.
 - *The cleaning blade* cleans excess toner off the drum after the print process has completed.



Figure 5-2: Identifying the toner cartridge found in laser printers.

- ◆ **Power supply:** The power supply in the printer is responsible for converting alternating current from the wall outlet into direct current that charges the primary corona wire and transfer corona wire as well as with other components of the printer.
- ◆ **Primary corona wire:** The primary corona wire applies the initial $-600V$ charge to the drum.
- ◆ **Transfer corona wire:** The transfer corona wire gives the paper a positive electrical charge that is used to attract the toner to the paper.

Each of the components of the laser printer is used to perform the print operation. The process that is used to perform the printer operation is known as the *laser printing process*.

The laser printing process

There are six phases to the laser printing process, and you are required to know them for the exam. Knowing the process is the basis for effective printer support and is essential for passing the A+ exam. The six phases of the laser printing process are conditioning, writing, developing, transferring, fusing, and cleaning. As you read the following sections, identify where each step occurs in the schematic in Figure 5-3.




Figure 5-3: Identifying the laser printing process.

1. Charging the drum (Conditioning)

When the printer receives a command from the computer's operating system to begin the print process, the primary corona wire applies a -600V charge to the *photosensitive drum*, also known as the *print drum*. This charge is one of the reasons why the printer requires a high-voltage power supply.

2. Exposing the drum (Writing)

After the drum has the -600V charge, a laser beam is used to hit areas on the drum to create the image that needs to be printed. In the areas on the drum that the laser touches, the charge changes from -600V to approximately -100V . Recognize that the areas exposed to the laser beam are more positively charged.

3. Developing the image (Developing)

After the image has been created on the photosensitive drum, the toner is used to develop the image on the drum. Alongside the print drum is a roller called the *developing roller*. The developing roller has a -600V charge, which attracts the toner from the toner reservoir to the developing roller.

Because the print drum and the developing roller are both charged with -600V (except for the areas of the print drum previously exposed to laser light), the toner from the developing roller is attracted to the -100V charged areas of the print drum. This entire concept is based on the “opposites attract” principle. Although both the drum and the roller are both negatively charged, -100V is more positive than -600V , so the toner on the -600V roller is attracted to -100V areas on the drum. Now that the print drum has toner on only the areas of -100V charge, the image is ready for transfer to paper.

4. Transferring the image (Transferring)

After the toner is on the print drum, the feed rollers (also known as the registration rollers) feed the paper into the printer and over the transfer corona wire. The *transfer corona wire*, also known as the *secondary corona wire*, applies a very strong positive charge of $+600\text{V}$ to the paper. The purpose of such a strong charge to the paper is to ensure that the toner will be attracted from the -100V areas of the drum to the paper. This, too, is based on the rule that opposites attract!

The paper continues to move through the assembly and passes over the drum to attract the toner from the drum to the paper.

5. Fusing the image (Fusing)

After the paper moves past the print drum and holds the toner, the paper then moves through the fusing rollers, which melt the toner to the paper. The fusing rollers are needed because the only thing holding the toner to the paper up to this point is a positive electric charge. During the fusing phase, the paper moves between a heated, Teflon-coated roller and a rubber roller, which melt the toner in place. The paper is then ejected from the printer.

6. Cleaning up the mess (Cleaning)

After the printing has completed, any excess toner that remains on the print drum needs to be cleaned off. That’s the purpose of this last phase — the cleaning phase. The cleaning phase of the laser printing process uses a cleaning blade that scrapes any leftover toner off the print drum and into a holding tray to prepare the drum for the next print operation.

The A+ Certification exams focus on laser printers when it comes to asking questions about printers. Be familiar with dot matrix and inkjet printers, but most of all be sure you are comfortable with laser printers and the laser printing process.



Lab 5-1 will help you identify parts of the laser printer. You will need a laser printer available to perform this lab. Lab 5-1 can be found in the Labs .pdf file in the Author directory of the CD-ROM.

Inkjet printers

Inkjet printers (see Figure 5-4) offer the next highest level of print quality and are relatively cheap compared to laser printers. Inkjet printers are great for home use or small office environments that don't have large print jobs.

Inkjet printers don't use toner like a laser printer; they use ink cartridges. The ink cartridge contains all the working elements needed to get an image from the computer onto a sheet of paper. It contains compartments of ink that are sealed with a metal plate to prevent the ink from running out. Each compartment has a tiny pinhole that is used to spray the ink out of the cartridge and onto the paper.




Figure 5-4: Looking at an inkjet printer.

When the printer receives the command from the computer to print an image, the printer starts the print process by applying an electrical charge to the heating elements that are in the ink reservoir. The charge heats the heating elements, which cause the ink to vaporize. The vaporized ink creates pressure that is forced out the pinhole, creating a tiny bubble that hits the paper.

Color inkjet printers are very popular today because of the increased popularity of digital cameras. Color inkjet printers may require two cartridges: one for black ink and one for the colors (cyan, yellow, and magenta). Most inkjet printers today have cartridges that bundle the black ink with the other colors. These cartridges are called *CYMK* (C for cyan, Y for yellow, M for magenta, and K for black).

Some printer manufacturers offer individual cartridges for each color. The benefit of these printers is that if you run out of one color you simply need to buy the cartridge that contains that color, not all of the colors.



When it's time to replace a spent ink cartridge, take the old cartridge with you to the store so that you know which cartridge type to buy. Some office supply stores can recycle your old cartridges, sometimes even offering a discount if you turn your old cartridges in. If you don't want to take the old cartridge with you to the store, make sure you know the make and model of the printer you're buying the cartridge for.

Dot matrix printers

Dot matrix printers are considered *impact* printers because they physically strike an inked ribbon with a metal pin to put characters on paper. A dot matrix printer (shown in Figure 5-5) fires off rows of pins that strike the ribbon in patterns to create the image or characters that need to be printed.



Figure 5-5:
Looking at a
dot matrix
printer.

Each pin, called a *solenoid*, is wrapped in a coiled wire that is held in place with a spring and small magnet. When a solenoid is needed to help create the image by striking the ribbon, an electrical charge is sent down the coil wire that surrounds the solenoid. The electrical charge around the wire causes the magnetic field from the magnet to be lost, resulting in the pin firing against the ribbon.

The solenoids are contained in the *print head*, which moves across the paper printing one line of dots at a time for the characters or image that needs to

be printed. Originally, dot-matrix printers used only nine pins in the print head. The 9-pin dot matrix printers were known as *draft-quality* printers and were later replaced by 17- and even 24-pin dot matrix printers. The quality of the 24-pin dot matrix was much better than that of the 9-pin because the greater number of dots creates a finer image.

Thermal printers

For the A+ exam, you simply need to focus on laser, dot matrix, and inkjet printers. But another type of printer you may encounter is a thermal printer. A *thermal printer* creates printouts on special paper by heating a stylus pen located on a print head. The pen then causes a chemical reaction on the special paper that is sensitive to heat.

Thermal printers are popular in restaurants for receipts because of the speed of the printer and how quiet it is.

Viewing Types of Printer Connections and Configurations

In this section, you find out how printers are connected to computer systems in order to communicate. There are a number of different ways to connect a printer to a computer — the following are the most popular methods:

- ◆ **Parallel connections**
- ◆ **Network connections**
- ◆ **Universal Serial Bus (USB)**
- ◆ **Serial**
- ◆ **Infrared**

Parallel connections

The most popular method for connecting a printer to a computer was, for many years, the parallel port. The end of the parallel cable that connects to the computer is a male DB-25 connector, while the end that connects to the printer is a male 36-pin Centronics connector. Figure 5-6 displays a parallel cable with the DB-25 and Centronics connector.

The maximum length of a parallel cable is 10 feet long — any longer and the cable runs the risk of *crosstalk* or *data skew*:

- ◆ **Crosstalk** is electrical interference from other equipment, fluorescent lights, and other cables.

Figure 5-6:
A parallel cable uses a 25-pin/36-pin connector on either end of the cable.



- ◆ **Data skew** is the concept of the signals that travel down the different wires in the parallel cable not traveling at the same speed and thus arriving at the destination at different times. This results in the data being unreadable at the opposite end.

Parallel connections deliver data 8 bits at a time at a speed of approximately 150 Kbps. Parallel printers and their cables should conform to the IEEE 1284 standard for parallel cables. This standard addresses parallel communication to and from the device attached to the cable.

Serial cable connections

Serial cables, which aren't very popular for printers, use either a 9-pin or 25-pin connector. Serial connections send data one bit at a time and are not susceptible to data skew. The maximum length of a serial cable is 25 feet long.

Network cable connections

Network-based printers have built-in network cards that allow the printer to connect directly to the network. The printer runs the TCP/IP protocol and is assigned an IP address so that it can participate on the network.

There are a couple of advantages to a network-based printer:

- ◆ **The network printer is available all the time.** Network-based printers are connected directly to the network, so they don't rely on a computer being on in order to communicate with the network.
- ◆ **Dedicated print servers are not required.** A network printer can be accessed from anywhere in the network, and you don't need to have a server in order to print to it. You may print to the network-based printer from the client computer directly.

Universal Serial Bus (USB)

Most printers today that are purchased for home or small-office use are USB printers, meaning that they connect to the computer via a USB port. Figure 5-7 shows a USB connection.



Figure 5-7: A USB connector is a popular method today to connect printers to computers.

USB has a number of benefits, including the fact that it is a Plug and Play technology — meaning that you can plug the device in without shutting down the system. USB 1.1 has a transfer rate of 12 Mbps, while USB 2.0 has a transfer rate of 480 Mbps.



The following are some key points to remember about USB:

- ◆ It is Plug and Play.
- ◆ You may connect 127 devices to a USB chain.
- ◆ You don't need to configure ports, IRQs, or DMA channels for each device.
- ◆ USB 2.0 has a transfer rate of 480 Mbps.



To find out more about USB ports, check out Book II, Chapter 1.

Lab 5-2 will give you the opportunity to practice connecting a USB print device to your computer. Lab 5-2 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Infrared

Connecting your printer to a system by using infrared technology lets your computer communicate wirelessly with the printer in much the same way that your TV remote lets you change channels without getting off the sofa. The infrared signal that is sent from the computer to the printer is carried as a beam of light, instructing the printer what to print.

In order to use an infrared printer, you need both an infrared transmitter/receiver connected to your computer and an infrared printer. Today, most computers have an infrared port (transmitter) built-in, especially on laptop systems. Desktop computers normally have a USB device that is installed to give an infrared transmitter.

There are three different types of infrared devices:

- ◆ **Reflective infrared:** The transmitters send the signal to a central unit, which then redirects the commands to the printer. This allows a number of users to print to a single printer at one time.
- ◆ **Line-of-sight infrared:** With line of sight, the printer's receiver must be in a direct line of site with the computer's transmitter. If there is a break in the line of site, communication is lost.
- ◆ **Scatter infrared:** Scatter infrared allows the signal to bounce off walls, ceilings, or people, all the way to the printer's receiver. The benefit is that you don't lose communication as you do with line of sight, but scatter infrared has limited range and transmission speeds.

FireWire and SCSI

You may also find printers that connect to a system via FireWire and SCSI connections. FireWire has two versions — one that runs at 400 Mbps and

one that runs at 800 Mbps. FireWire 800, also known as FireWire 1394b, is a fairly recent update to the FireWire technology.

SCSI is another common type of printer device that you may come across with Macintosh systems! Most PC-based systems don't use SCSI printers, although it is possible. For more information on SCSI devices, check out Book II, Chapter 5.

Upgrading Printer Hardware

Printers aren't nearly as upgradeable as computers. There are only a few items that may need upgrading on a printer:

- ◆ **Printer memory**
- ◆ **Disk space on larger printers**
- ◆ **Firmware updates**

I discuss these upgrades in the following sections.

Upgrading memory

A printer uses memory to store the information that needs to be printed, so the more memory a printer has, the larger the print documents it can store. If you ever receive an insufficient memory message when printing a large document, your printer probably doesn't have enough memory. You'll need to buy and install some more memory for the printer if you have this problem.

Adding a disk drive

On larger printers, you can add hard disk space, typically through a PCMCIA slot on the printer. The additional hard drive can store information used by the printer; for example, you may want to install some fonts on the printer, or you could have a print queue stored on the actual print device and use the hard drive to store the queued documents.

Upgrading the firmware

You may also need to update the firmware on the printer. The firmware dictates the capabilities of the printer and is similar to the BIOS of a computer. Actually, it is so similar that you upgrade the firmware in a similar way as upgrading a BIOS. You need to download the firmware update from the manufacturer's Web site. This update is usually in the form of a self-extracting executable. After you have downloaded the update, you double-click it to run the update.



Make sure the printer is connected and powered on when the update runs because the update program needs to rewrite the printer firmware within the physical printer.

Installing a Printer in Windows

After you have connected the printer to the computer, typically with a parallel or USB connection, you then need to install the “printer” into Windows. Windows uses the term *printer* to describe the software interface (the printer driver) used to communicate with the print device. This software interface is represented as the icon that displays in the Printer folders — the icon that you right-click to change the settings for the corresponding print device. Microsoft says that the *printer* is the icon found in Windows, while the *print device* is the piece of hardware that is connected to the computer. I use these terms when discussing how you manage a printer in Windows.

Installing a printer

When you connect a USB print device to a Windows computer, Plug and Play kicks in and detects the hardware. If Windows has a driver for the print device, it will load that driver automatically; if it doesn't, it will prompt you for the driver disk. After you supply the driver disk, the printer is installed, and you are off to the races!

For the following example, assume that you have a parallel print device that is connected to LPT1. You need to install a printer in Windows that connects to that device so that you can print. To install the printer in Windows XP, follow these steps:

1. Choose Start⇨Printers and Faxes.

The Printers and Faxes window appears, and you will notice an Add a Printer icon.

2. Click the Add a Printer link to install a printer driver.

The Add Printer Wizard magically appears.

3. Select whether you are installing a local printer or a network printer, as shown in Figure 5-8.

- A *local printer* is managed by the system. A local printer is not necessarily physically connected to the system, but it is a resource of the local system — meaning that others will connect to the computer to get to the resource.
- A *network printer* does not exist on your system but lives somewhere else on the network. You still need to install the printer on your computer so you have the driver, but it is not your resource, so it is a network printer.

Figure 5-8:
Installing a
local printer
in Windows
XP.



4. Also on the screen shown in Figure 5-8, specify whether you want Windows to try to detect the type of printer that is connected to your system.

You will select the printer you are installing, so turn off that option and then click Next.

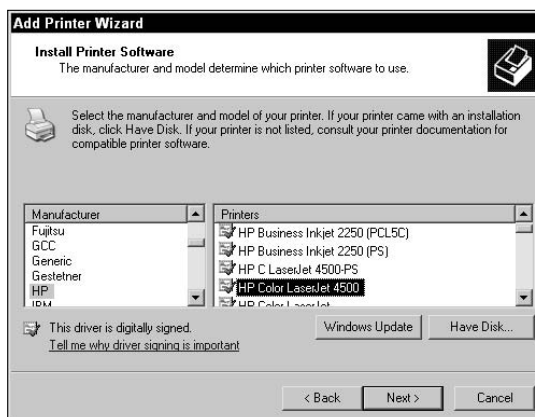
You are then asked what port the printer is connected to.

5. Make sure that LPT1 is selected and then click Next.

6. Select the type of printer you wish to install (as shown in Figure 5-9) and click Next.

Choose the manufacturer on the left side and then the model of the printer on the right side. If your printer does not appear in the list, you can click the Have Disk button to provide the driver CD that came with the printer.

Figure 5-9:
Choosing the manu-
facturer and
model of the
printer.



7. Click **Next** to accept the default name of the printer.
8. Indicate whether you would like to share the printer and then click **Next**.

Sharing the printer allows other users on the network to connect and print to it from their computers. For now, don't share the printer.
9. If you want to print a test page to verify that the printer you are installing works, choose to print a test page.
10. Click **Next** to go to the summary screen in the **Add Printer wizard**.

A summary page appears, displaying information about your newly installed printer.
11. Click **Finish** to complete the printer installation.

Configuring a printer

After the printer has been installed, you may want to configure it. To configure the printer, choose **Start**⇨**Printers and Faxes**, right-click the printer, and choose **Properties**. The printer's **Properties** dialog box opens. The following is a list of popular printer settings that may need to be changed depending on how you want to use the printer:

- ◆ **Priority:** If a group of users needs to have a higher priority on a print device over others on the network, you could install two printers to refer to the same print device. After you have both printers installed, set the priority of one of the printers higher than the other — anyone's print jobs sent to that printer will have a higher priority on the print device. You can set the priority of a printer on the **Advanced** tab of the printer's **Properties** dialog box (shown in Figure 5-10).
- ◆ **Schedule:** If you have a group of users that should be printing only at certain hours, you could set the schedule of the printer so that it can print only during those times. The user can send print jobs to the printer and the printer will queue the job, or store it, until the scheduled time. At the scheduled time, it will then print the job. The **schedule** option can also be seen in Figure 5-10.
- ◆ **Spool settings:** In Windows, when a user chooses to print a document, the print job creates a temporary file on the disk for the document being printed. After the file is stored on disk, the user gets control of the application and the system sends the temporary file to the print device to be printed. While the temporary file is sent to the print device, the user gets to use his or her program again. This is sometimes referred to as **background printing** — where the user thinks the print job has been sent, but in actual fact the print job is being sent while the user continues using the computer. The purpose of this process, called *spooling*, is

so that users can continue using the computer instead of waiting for the 20 pages to actually get sent to the print device. Because temporary files are stored on the disk when spooling is enabled (which is the default with all Windows printers), if you don't have the hard disk space on your computer to store the temporary files, you won't be able to print. If you get a spooling error, you may turn off spooling by choosing Print Directly to the Printer on the Advanced tab. Again, you can see this option if you look back at Figure 5-10.

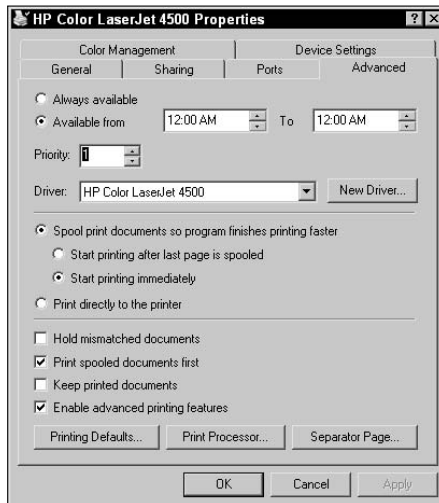


Figure 5-10: The priority setting and the scheduling option on a printer.

- ◆ **Driver:** If the driver goes corrupt for your printer, you may want to change the driver. To change the driver for a printer, go to the Advanced tab of the printer's Properties dialog box and click the New Driver button. You will then be asked for the make and model of the printer.
- ◆ **Print Test Page:** If you run into trouble with a printer, you may want to print a test page by going to the Properties dialog box of a printer and choosing the General tab. On the General tab, click Print Test Page. Printing a test page will help you determine if you are having a problem with the printer or a problem with an application. For example, if you cannot print from Microsoft Word but you can print a test page then there is a setting in Microsoft Word that is causing the printing problems.
- ◆ **Sharing:** If you're in a networked environment, you can share the printer so that other users on the network can print to it from their computers. To share the printer, go to the properties of the printer and click the Sharing tab. Select the Share This Printer option and give the printer a share name, as shown in Figure 5-11.

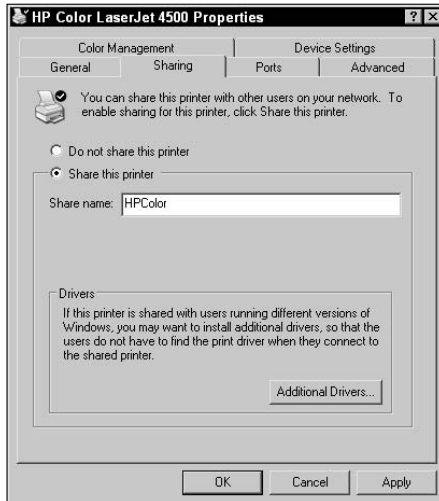


Figure 5-11:
Sharing a
printer.

Connecting to the shared printer

When you have shared a printer, users need to install a printer on their systems that points to the shared printer on the network. Users who install a printer that refers to a shared printer on the network are installing what is called a network printer. A *network printer* refers to a shared printer on the network by what is known as a *universal naming convention (UNC)* path. A UNC path always has the syntax of

```
\\computername\sharename
```

The `computername` parameter is the name of the computer that has the printer shared, whereas the `sharename` parameter is the name the printer has been shared as. For example, assume you have a system with a computer name of `WORKSTATION1`, which has a printer shared as `HP`. To connect to this printer, you would type `\\workstation1\hp` from the Run command or when installing the network printer.

To install a network printer and connect to a shared printer, follow these steps:

- 1. Choose Start⇨Printers and Faxes.**

- 2. Click the Add Printer link.**

The Add Printer Wizard appears.

- 3. Click Next on the Welcome screen.**

You are asked if you want to install a local printer or network printer.

4. Select **Network Printer** and click **Next**.
5. Select the **Connect to This Printer** option and then type the UNC path of the shared printer on the network (shown in Figure 5-12).



Figure 5-12: Connecting to a shared printer.

6. Click **Next**.
7. Click **Finish**.



Lab 5-3 will give you the opportunity to practice installing a printer in Windows XP. Lab 5-3 can be found in the `Labs.pdf` file in the `Author` directory of the CD-ROM.

Troubleshooting Printer Problems

Printers are wonderful devices that we depend on day in and day out as we prepare, conduct, and report on our business. But one of the major downfalls of a world that depends on printers is that they never seem to work!

You discover a number of general troubleshooting guidelines in the sections that follow to help you troubleshoot printer problems. You also find out about a few specific problems with each type of printer.

Check the simple stuff first

The first thing you want to do when you cannot print is to verify the simple stuff. Is the printer connected, powered on, and online? Visually check the printer itself and make sure that you can see the online lamp that indicates that the printer is ready for use.

Paper jams

Another type of problem that occurs a lot is paper jams. If you experience a number of paper jams over and over again, then you may want to verify that you are using the correct type of paper for the printer. The best thing to do here is to check the documentation of the printer and verify that you are using the correct size paper. If you are using paper that is too thick for your printer, it may jam up a lot.

Garbled or corrupted output

If you experience corrupted output of any type, it is possible that the driver has gone corrupt or that you have the wrong driver installed. One of the best ways to find out if you have a bad driver or if the application you are printing from is causing the problem is to print a test page.

To print a test page in Windows XP, follow these steps:

- 1. Choose Start⇨Printers and Faxes.**
- 2. Right-click the troublesome printer and choose Properties.**
- 3. On the General tab, click the Print Test Page button.**
- 4. Click OK if the page printed properly or click the Troubleshoot button if the page did not print correctly. A troubleshooting Web page will appear, giving you a number of tips on how to troubleshoot the problem.**

If you print a test page and the printout is still garbled, the driver has likely gone corrupt, and you need to install a new driver by going to the properties of the printer and choosing the Advanced tab. On the Advanced tab, click the New Driver button to install a new driver.

Spots or smudging on the printout

In general, spots or constant smudging on printouts are good indications that you're using the wrong type of paper for your printer. Again, check the documentation for the printer to find out what type of paper you should be using.

Slow printing

If you notice that printing is overly slow, then you should verify that spooling is enabled. Spooling is enabled by default, but it may have been switched to Print Directly to Printer when you were troubleshooting or configuring the printer. Check the properties of the printer and ensure that spooling is enabled.

Spooling service problems

In Windows, the *print spooler service* is responsible for managing the printing environment. If you notice that a print job is hung in the print queue and will not print or you cannot delete it, then you may have a corrupt queue.

When this happens, you need to stop and restart the print spooler service in Windows. Stopping the print spooler service deletes all print jobs and essentially “reboots” your printing environment for you. Figure 5-13 shows the print spooler service being restarted, which you can do from the Services console. The Services console is found in Start | Control Panel | Performance and Maintenance | Administrative Tools.




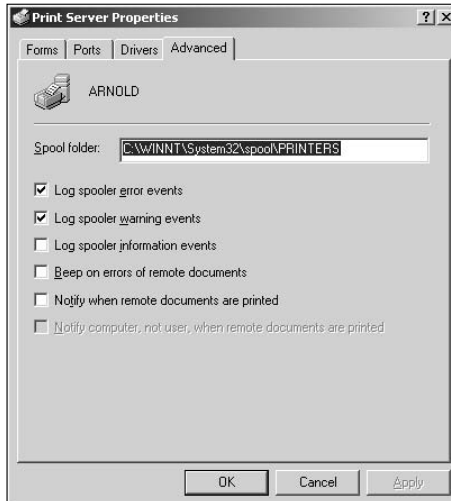
Figure 5-13: Editing the spooler services properties on a Windows 2000/XP system.

If you have a number of problems printing in Windows and you have determined that the problem is not hardware-related, you may need to move the print spooler folder. By default, the print job is spooled to the hard drive at the `%systemroot%\system32\spool\printer` directory — the `%systemroot%` variable is typically the Windows folder on drive C.

If you are running out of space on drive C, you may want to change the partition for the spool directory. To change the default spool folder in Windows, follow these steps:

1. Choose Start⇨Printers and Faxes.
2. Choose File⇨Server Properties.
3. Click the Advanced tab, shown in Figure 5-14.

Figure 5-14:
Changing
the spooling
folder on a
Windows
system.



4. Enter the path of the new Spool folder.


If storage or drive speed is your motivation, this should be on a different partition or drive.



Dot matrix problems

This section identifies some popular problems that occur with dot matrix printers. Be sure to review these problems and the possible causes before taking the A+ exam:

- ◆ **Faint printing:** If you experience faint printing with your dot matrix printer, the print ribbon is simply worn out. You need to replace the ribbon.
- ◆ **No printing:** If your dot matrix printer simply doesn't print, a print head cable might be disconnected, or the print head might have torn through the ribbon. In these cases you will need to connect the print head cable or replace the ribbon.
- ◆ **Paper jamming:** Again, a lot of paper jams is a great indication that you have the wrong type of paper or the wrong size paper. Check the documentation for your printer.
- ◆ **Line across the page:** If your dot matrix printer prints a line all the way across the page, you might have a pin in the print head that is stuck out. You may be able to loosen the stuck pin or you may need a new print head.



Inkjet problems

This section identifies common problems that can occur with inkjet printers during day-to-day activity. Again, review these before taking your A+ exam:

- ◆ **Paper jam:** A paper jam indicates that you may have the wrong type or size of paper. If you have the correct size paper, the feeder wheels might be dirty — they are responsible for moving the paper through the printer.
- ◆ **Poor print quality:** Having a poor-quality printout could indicate that the ink needs to be replaced or that you have the wrong type of paper. Change the ink or check the documentation to ensure you are using the correct type of paper.
- ◆ **Fading print:** If the print from your inkjet printer is fading, that could be an indication that you need to change the ink.

Laser printer problems

The following common problems may occur with a laser printer:

- ◆ **Faint print:** If you find that the print is getting faint in your laser printouts, then you most likely need to replace the toner.
- ◆ **Paper jam:** If you find you get a lot of paper jams, you need to verify that you are using paper of the correct size and thickness. If you're sure you have the right paper, you could also have a problem with misaligned rollers.
- ◆ **White stripes:** If you find you have white stripes throughout the printout, then your transfer corona wire is most likely the problem — there may be a problem getting the charge to the paper, resulting in toner missing on the page in areas.
- ◆ **Blank page:** If you have a blank page for output, you know there is nothing wrong with the feeder mechanisms because the paper is moving through the printer. A blank page indicates a problem getting the toner to the paper, so there is something wrong with the corona wires or you have no toner.
- ◆ **Vertical line:** A vertical line on the printout typically indicates that there is a scratch in the print drum — you should replace the toner cartridge because it contains the print drum. If you replace the toner cartridge and the problem still exists, then there could be a problem with the laser.
- ◆ **Smearred pages:** If toner smears off after the page has printed, something is wrong with the fuser rollers. You may need to replace the fuser rollers. Be cautious when working with the fusing components because they are very hot.

Understanding Safety and Preventative Maintenance

In this section, I show you some common safety and preventative maintenance practices for printers.



For the exam, be sure to review the safety points in Book I, Chapter 3, along with the safety and preventative maintenance points presented here.

Safety precautions

The first point to make about working with printers is to be sure you turn the printer off and unplug it before doing any maintenance on it. Also be sure to give the printer time to cool down as there are many parts in the printer that can get very hot — especially with laser printers.



If you are servicing a laser printer, be sure to give the printer time to cool down because the fuser rollers get very hot. These rollers are hot enough to melt toner to the paper, so they can cause severe burns to your skin — be careful!

You also want to be cautious when working around toner — it can get very messy. Be sure not to get toner in your eyes or on your skin. If you have spilled toner, get a vendor-approved vacuum to suck it all up.

Preventative maintenance

Preventative maintenance is the process of taking care of the printer to help reduce downtime. Think of your relationship with your printer like any other relationship — if you treat the printer well, it will stick around; if you don't treat it well, it'll leave you when you least expect it! Care for your printer by cleaning it every once in a while — at least once a year (make it an anniversary thing!).

To clean your laser printer, follow these general steps or consult the manufacturer's documentation for exact methods of cleaning your printer make and model. Some general steps to follow to clean the laser printer are:

- 1. Turn off and unplug the printer.**
- 2. Clean the outside of the printer with a lint-free cloth and a tiny bit of isopropyl alcohol.**
- 3. Remove the paper tray and paper. Clean the tray, removing any dust or paper particles. Clean the area inside the printer where the tray resides. Be sure to clean the feeder rollers — they get pretty dirty.**
- 4. Open the printer and remove the toner.**

5. Clean the inside of the printer and again don't be afraid to clean the rollers.
6. Put the toner and the paper tray back in and print a test page.



Lab 5-4 allows you the opportunity to review the steps to cleaning an ink cartridge on a inkjet printer. Lab 5-4 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Getting an A+

In this chapter, you find out about three types of printers — dot matrix, inkjet, and laser printers. You discover how each printer prints and find out about the six phases of the laser printing process. The following are some key points to remember about printers when preparing for the exam:

- ◆ Laser printers use light, electromagnetism, and heat to create a printout.
- ◆ Inkjet printers use a heating element to vaporize ink and spray it through a pinhole to make characters on a page.
- ◆ Dot matrix printers use solenoids to strike a ribbon, creating characters out of patterns of printed dots.
- ◆ You may use a parallel connection, USB, or network connection to send the print job from the computer to the printer. USB is popular for home and small-office use, while larger companies use a printer with a built-in network card.
- ◆ The six phases to the laser printing process are charging (conditioning), writing, developing, transferring, fusing, and cleaning.
- ◆ You may install a printer in Windows by running the Add Printer Wizard. After the printer is installed, you can configure a number of settings, such as spooling, scheduling, and priority.

Prep Test

- 1** What are the two types of feeder mechanisms used in printers?
 - A Continuous tractor feed
 - B Continuous form feed
 - C Friction feed
 - D Injected

- 2** Which of the following devices use friction feeder mechanisms? (Choose two.)
 - A Laser printers
 - B Fax machines
 - C Dot matrix printers
 - D A printer that prints hundreds of payroll checks

- 3** A positive charge is applied to the paper by which laser printer component?
 - A The laser
 - B The print drum
 - C The registration rollers
 - D The transfer corona wire

- 4** What is the purpose of the negative charge on the print drum of a laser printer?
 - A To attract the toner to every area of the drum
 - B To attract the toner to the areas of the drum that have a stronger negative charge
 - C To attract the toner to the areas of the drum that have a weaker negative charge
 - D To attract the positively charged paper to the print drum

- 5** Why are both the developer roller and the print drum charged with $-600V$ in a laser printer?
 - A So the paper is attracted to neither
 - B So the toner is attracted to neither
 - C So the toner creates a fusion cloud between the two rollers and the paper
 - D So the toner is attracted only to weakly charged areas of the print drum

- 6** The paper is charged with which voltage charge by the secondary corona wire?
- A +600V
 - B -600V
 - C -100V
 - D +100V
- 7** What method does an inkjet printer use to print a page?
- A One dot at a time to form a character
 - B Spray-painting a character
 - C Striking an inked ribbon
 - D Dropping ink onto the paper
- 8** With an inkjet printer, what causes the ink to vaporize?
- A Electrical charge
 - B A heating element within the ink cartridge
 - C A solenoid in each chamber of the ink cartridge
 - D Drying of the ink when the cartridge has not been used for some time
- 9** A vertical line on every page of a printout from a laser printer indicates what?
- A Fleck of toner
 - B Scratch on the print drum
 - C Light leakage
 - D Fingerprint
- 10** What would you do if the print job refused to leave the print queue in Windows 2000/XP?
- A Delete the printer and reinstall it
 - B Restart the computer
 - C Turn the printer off and then on
 - D Stop and restart the spooler service

Answers

- 1** B, C. Continuous form feed and friction feed are the two types of paper feeder mechanisms. Dot matrix printers use continuous form feed, and laser printers use friction feed. *See "Learning about Paper Feeder Mechanisms."*
- 2** A, B. Friction feeders use rollers that apply pressure and friction to the top sheet of paper in the paper tray to move it into the print device. Laser printers and fax machines use friction feed. *Review "Friction feeders."*
- 3** D. The transfer corona, also known as the secondary corona wire, charges the paper with a +600V charge. *Check out "The laser printing process."*
- 4** C. The print drum is charged with an initial -600V charge. The laser writes the image onto the print drum by weakening the charge to -100V in areas that create the image. This causes the toner to attract to the -100V areas because the toner is approximately -600V — and opposites attract! *Peruse "The laser printing process."*
- 5** D. The developer roller and the print drum are both equally charged so that they are repelled from one another. The toner on the developer roller will be attracted to neutralized areas of the print drum caused by the laser. *Take a look at "The laser printing process."*
- 6** A. The transfer corona charges the paper to +600V so the toner on the print drum is strongly attracted to it. *Peek at "The laser printing process."*
- 7** B. An inkjet printer sprays the ink onto the paper. *Look over "Inkjet printers."*
- 8** B. Within the ink cartridge, a heating element vaporizes the ink, forcing it to spray a drop out the pinhole in the nozzle. *Study "Inkjet printers."*
- 9** B. A scratch on the print drum can cause the vertical line to appear in the print-out. *Refer to "Laser printer problems."*
- 10** D. You would stop and start the print spooler service, deleting everything from the queue. Stopping and starting the service is a way of rebooting the printing environment. *Examine "Spooling service problems."*

Chapter 6: Working with Multimedia Devices

Exam Objectives

- ✓ **Understanding scanners**
- ✓ **Understanding digital cameras**
- ✓ **Working with other multimedia devices**

It's amazing to see an old uncle at a birthday party who, just a few years back, hadn't even touched a computer, but who now snap shots with his digital camera and drones on about how simple it is to upload the pictures to his computer and then burn them to CD-ROM! When I try to think of examples of how computers have entered our daily lives, this is the most promising example I can think of.

In this chapter, you find out about popular multimedia devices used in everyday life. In the following pages, I discuss a number of multimedia devices, starting with digital cameras and scanners — the two most popular types of computer-related multimedia devices. I also introduce a few key points about microphones and MIDI technology, which are very important to the many people who will record music through their computers.

Understanding Scanners

In this section, you find out about different types of scanners. I also show you how to scan a photo in Windows XP on a scanner and how to save the photo as a file on the computer.

A few years back, scanners were very popular because they were one of the only ways to get your photos into a digital format. Maybe you wanted to use the computer to alter the photo or to send it to a relative via e-mail; to do any of these actions, you had to use a scanner, which would copy the photo as a digital image that could be saved on your computer.

Many different types of scanners have been manufactured over the years, and you need to be familiar with the different types of scanners for the A+ Certification exam. They are:

- ◆ **Flatbed:** The flatbed scanner is the most popular scanner today. With a flatbed scanner, you lift the cover of the scanner and place the photo on the flat sheet of glass (called a *platen*). A scan head then moves beneath the picture, scanning the image by using a piece of technology called a *charge-coupled device (CCD)*. Figure 6-1 shows a flatbed scanner.
- ◆ **Handheld:** A handheld scanner is great in environments where it is impossible to place the object to be scanned on a flatbed scanner. To use a handheld scanner, you move the scanner over the object that you want to scan.
- ◆ **Sheetfeed:** A sheetfeed scanner acts very much like a fax machine in the sense that the scan head is stationary. You insert the paper object to be scanned, and the feeder passes the paper over the scan head. This scanner type looks very similar to a portable printer and is also the scanner type typically found in multifunctional printers.
- ◆ **Drum:** A drum scanner uses a technology called a *Photomultiplier Tube (PMT)*, which involves wrapping the photo around a glass tube. The drum releases light that is converted to electrical data to become the image file.



You need to be familiar with the different types of scanners for the A+ Certification exam.




Figure 6-1: A flatbed scanner uses a scan head that moves under the object being scanned.

Scanning process

Most scanners use a *charge-coupled device (CCD)*, which is responsible for capturing the photons that are created by the light being exposed to the picture and then converting those photons to electrons, which are a form of electrical data. The image is then created from the electrical data.

The scanner has a scan head that contains the CCD along with a number of mirrors and lenses. The mirrors reflect the light from the lamp onto the photo being scanned. That light then passes through the lenses to the CCD to create the image.

Connecting a scanner

You can use a number of different types of ports to connect a scanner to the computer, depending on what type of connection your scanner wants to use. The popular ports used by a scanner are as follows:

- ◆ **Parallel port:** Older scanners connected to the system by using the parallel port, which might have caused a little problem because the printer was already using the printer port! Don't despair — if you look at this type of scanner, you'll notice that it has two connectors on it. You can connect the scanner to the computer and then, with another cable, connect the printer to the scanner, creating a *daisy chain*. In this scenario, you need to check your CMOS settings to ensure that the parallel port mode is set to EPP (enhanced parallel port), which supports daisy chaining.

Check out Book II, Chapter 4, if you need to review how to adjust your computer's CMOS settings.

- ◆ **SCSI:** Another popular form of scanner connection a few years back was the SCSI (Small Computer Systems Interface) connection. To connect a SCSI scanner to a computer, you need to ensure that you have an SCSI host adapter, that you assign a unique ID to the scanner, and that you make sure you terminate the SCSI bus.
- ◆ **USB:** The most popular method used to connect a scanner to a computer these days is the USB (Universal Serial Bus). With USB, you simply connect the scanner to the computer, and Plug and Play kicks in and either loads the driver or prompts you for it.

Scanning an image

In this section, I show you how to scan an image by using Windows XP. For this discussion, I assume that you have already connected the scanner and have loaded the driver for the device.

Follow these steps to scan a picture in Windows XP:

1. Choose Start→Control Panel→Printers and Other Hardware→Scanners and Cameras.

You should see your scanner in the Scanners and Cameras window.

2. To scan an image, select your scanner and click Get Pictures on the Imaging Tasks pane on the left side of the screen.

This launches the Scanners and Cameras Wizard, which copies pictures to your system.

3. Click Next on the wizard's welcome screen.

4. Indicate that you wish to obtain a color picture and then click Next.

5. Type the name for this group of pictures, as shown in Figure 6-2.

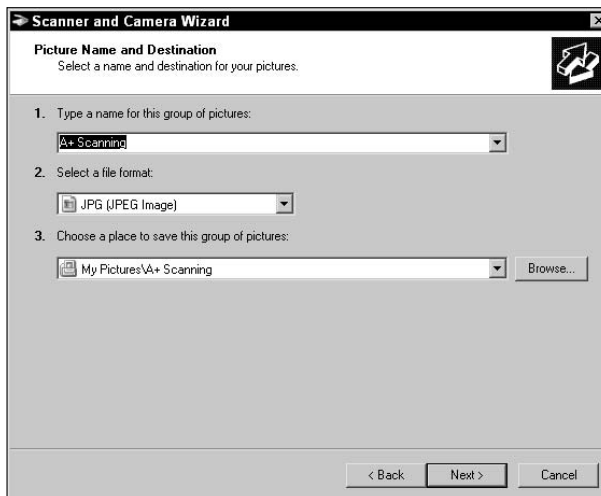


Figure 6-2:
Scanning a
photo in
Windows
XP.

The name you give to the group of pictures will be the name of a folder in My Documents that holds all the photos you scan. It will also be the first part of the name of each image (followed by an incremental number).

6. Select which graphics file format you would like to save your pictures in.

Graphics file formats are beyond the scope of this book, but here's a guideline: As a general rule, if you will use the image only for on-screen viewing, go with JPEG. If you hope to eventually print the image on paper, save it as a TIFF. The TIFF file will be larger, but it will include more detail than a JPEG file.

7. Click Next, and your scanner starts doing its thing.
8. When the scanning is complete, click Next and then click Finish.

A new folder will appear with your newly scanned image(s)!

Using Digital and Web Cameras

Just a few years ago, taking pictures involved buying film for your camera, snapping some pictures, and then taking the exposed film to a photo lab to be developed — often only to find out that of the 24 pictures you took, only four of them were any good!

Photography has changed dramatically since then. Today, you use a digital camera, which stores the picture in memory on the camera, not on film. You can also navigate through the pictures by using the built-in viewer on the camera and then delete the ones you don't like — then print only the ones you want developed!

The following sections introduce you to digital cameras, digital video cameras, and Web cameras, also know as *Webcams*.

Digital cameras

A digital camera is different from a conventional camera because, instead of storing the image on film, a digital camera stores the image as a digital file on the camera. This file is then uploaded to a computer and can be used in any application on the computer, such as e-mail, presentation software, or a photo editor.

The digital camera works the same way that a scanner does — it captures light and uses CCD to convert the captured light into an electrical charge that is used to create the pixels in the image. The light is captured through the lenses on the digital camera and is then passed through the CCD so that the CCD can do its job of converting the light into data.

When shopping for a digital camera, you want to be sure you get the best camera for your dollar. The quality of a digital camera is measured in its resolution and zoom levels.

Camera resolution

Camera resolution is measured in pixels, just like monitor resolution. The more pixels recorded in the photos being taken, the more detail that is displayed in the picture. The resolution also affects how large the picture can become before becoming grainy. A camera advertises its resolution by indicating how many pixels it supports, which is typically measured in

megapixels (or millions of pixels). The higher the number of pixels supported, the better quality image you will get out of your camera. Here are what some common camera resolutions translate into:

- ◆ **1 megapixel:** Supports an image resolution of 1216 x 912.
- ◆ **2 megapixel:** Supports an image resolution of 1600 x 1200.
- ◆ **4 megapixel:** Supports an image resolution of 2240 x 1680.
- ◆ **5 megapixel:** Supports an image resolution of 2560 x 1920.
- ◆ **6 megapixel:** Supports an image resolution of 2816 x 2112.
- ◆ **8 megapixel:** Supports an image resolution of 3264 x 2448.
- ◆ **11 megapixel:** Supports an image resolution of 4064 x 2704.

By way of comparison, an 11-megapixel camera is professional-quality. At the time of this writing, the typical consumer-grade digital camera has a resolution of between 4 and 6 megapixels.

Most digital cameras let you specify which resolution you would like to take photos in. For example, on my 3-megapixel camera, I can specify that I want to take pictures with a resolution of 2 megapixels, which creates a smaller image file but also creates a lower-resolution image. But if I'm happy with the quality of the 2-megapixel image (it's fine for eBay!), because they are smaller files, I will be able to store more of them on the camera.

Zoom

Digital cameras can have two types of zoom, optical and digital, both of which let you zoom in on your subject when taking a photo. In the following list, optical and digital are compared against what is known as *fixed focus*:

- ◆ **Fixed focus:** A fixed focus camera does not support any type of zooming features. What the camera sees is what the camera takes as a picture. To zoom in, you need to physically get closer to the object. This is the zoom type (actually a lack of zoom) used on disposable cameras.
- ◆ **Optical zoom:** Optical zoom is the “good” zoom type that uses lenses on the camera to change the focus and zoom in on the object. This zoom type does not alter the image in any way — it simply zooms in and takes the picture.
- ◆ **Digital zoom:** Digital zoom is what I like to refer to as a *virtual zoom*. It takes the picture and then goes to the center of the image and magnifies the subject. This is the same as you taking an original photo into a photo editor and magnifying the image and then cutting to the part you like. This typically never turns out right because, as you magnify the image, it becomes grainy.

Uploading pictures to the computer

To upload pictures from your digital camera in Windows XP, follow these steps:

1. Plug the camera into the computer's USB port and turn the camera on.

Windows XP detects the camera, loads the driver for the camera, and then asks which program you want to use to copy your photos to the system.

2. Choose Microsoft Scanner and Camera Wizard and then click OK.

3. When the Camera Wizard displays its introduction screen, click Next.

You are presented with a list of photos that can be uploaded.

4. Make sure that a check mark appears beside each photo you want to upload (as shown in Figure 6-3) and then click Next.

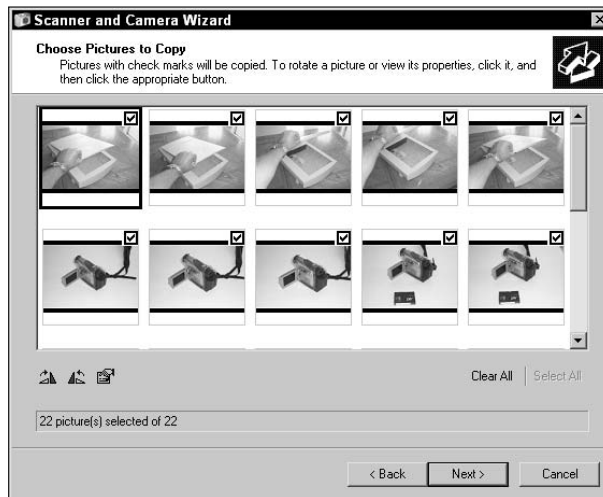


Figure 6-3: Uploading photos from a digital camera to your computer with Windows XP.

The Wizard asks you to name the group of photos you want to upload.

5. Type a name for this grouping and then click Next.

The photos are then copied to your system and placed in the My Documents\My Pictures folder.

6. When the photos have been copied, click Next and then click Finish.

Web and digital video cameras

In this section, I introduce you to Web cameras, digital video cameras, and the concept of a CODEC — all of which are important to understand when supporting video devices or video applications.

Web cameras

Web cameras (also called *webcams*) can capture video and display it on the computer screen, or they can be used to record a video and store it in a digital format on the computer. Webcams are popular with conferencing applications, which allow you to have a live conversation while seeing the person you are talking to, over the Internet. An example of an application that can use a webcam is MSN Messenger (see Figure 6-4).

Figure 6-4:
MSN
Messenger
lets you
(informally)
chat or
(formally)
video
conference.



The webcam is typically shaped like an eyeball with a lens on the front and connects to the system's USB port. Most webcams have a stand that holds the camera and allows it to sit on the monitor or desk. There are also laptop versions of webcams that clip to the frame around the laptop's LCD. Figure 6-5 shows a webcam.

Digital video cameras

Just as digital cameras have replaced conventional methods of taking snapshots, recording video with a digital video camera has become popular in this day and age. The digital video camera can record directly to the computer and store the video as a file on the computer, or it can store the video to memory on the camera. It can also store the information to a digital video tape that can then be played at a later time and recorded on the computer as a file.

A digital video camera typically connects to the computer via a USB 2.0 or IEEE 1394 (FireWire) connection, which is currently the more popular choice.



Figure 6-5: A webcam captures live video and sends it to anyone you are conferencing with.

Understanding CODECs

The problem with digital video is that, with all the information that needs to be recorded, the files can quickly become huge. Digital video file sizes are kept to a minimum by compressing the video using special *COmpression/DECompression (CODEC)* software. There are a number of different types of video formats and CODECs, the most popular of which are in the following list:

- ◆ **DV:** Digital video cameras have their own proprietary format to store the video, known as the DV format. The DV format typically has no loss in quality.
- ◆ **DivX:** DivX is one of the most popular CODEC types today. It supports high-quality video with high compression ratios. A video using the DivX CODEC uses only about 15 percent of the disk space that other CODEC types do.
- ◆ **MJPEG:** MJPEG is based on the JPEG compression standard used to compress still images. This compression type is optimized for transferring video to and from tape but is not suited for Internet applications because special hardware is required to play the video back.
- ◆ **MPEG1:** MPEG1, developed by the *Moving Picture Experts Group*, has been used heavily for creating videos to be distributed over the Internet because of its small file sizes and the fact that only the CODEC is needed to play the video back.
- ◆ **MPEG2 and MPEG4:** The MPEG1 video compression has been improved again and again to create the MPEG2 and MPEG4 compression standards, which have better compression rates and better resolution. For example, the MPEG2 standard offers four times the resolution of MPEG1.



The important point to remember with CODECs is that you must have the same CODEC installed on your computer that was used to record the video in the first place. Otherwise, the computer won't know how to decompress the video. If you don't have the correct CODEC, you won't be able to play the video.

Looking at Other Multimedia Devices

Two other multimedia devices that are popular in today's systems are MIDI devices and microphones. Sound cards typically have a microphone port (also known as a *mic port*) and a MIDI port.

Microphones

The microphone port on the computer is indicated with a little picture of a microphone. You can use the microphone port to plug a standard microphone into the computer. If you have a microphone whose cable will not connect in the mic port, then you need to head to your local electronics store and find an adapter. After you plug the microphone in, you can use it as a recording source as long as you have specified it as a recording source in Windows.

To set the microphone as a recording source in Windows XP, follow these instructions:

1. Choose Start⇨All Programs⇨Accessories⇨Entertainment⇨Volume Control.

Here you can set the levels for each type of input and output device. For example, when recording, you want to make sure that the level on the microphone is not too high or the recording will sound distorted.

2. To ensure that the microphone is set as a recording device, choose Options⇨Properties in the Volume Control dialog box.

The Properties dialog box appears. Notice that Playback is selected, which means that you are seeing and can select which devices will play back sound.

3. Select the Recording radio button and verify that the Microphone is selected, as shown in Figure 6-6.

4. Click OK and then close the Volume Control dialog box.

MIDI

Musical Instrument Digital Interface (MIDI) is a standard developed to allow musical instruments to communicate with one another or with a computer. MIDI is a common language that all MIDI-capable instruments can speak. In

order to use MIDI, you must have a MIDI-capable instrument — for example, a keyboard with a MIDI port on it — and a system with a MIDI port. You don't need to worry about the details of MIDI for the A+ Certification exams.

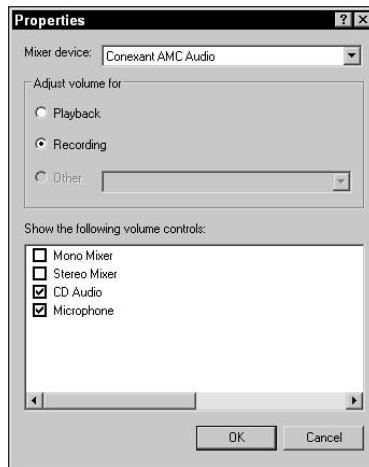


Figure 6-6:
Ensuring
that the
microphone
is set as a
recordable
device.

Getting an A+

This chapter introduces you to some of the different devices that are used when working with multimedia and multimedia applications. Here are some key points to remember about multimedia devices:

- ◆ A flatbed scanner is the popular type of scanner and uses a technology called a *charge-coupled device (CCD)*. The scan head on the scanner moves under the image, exposing the image to light that is captured and then converted into electrical data by the CCD, which creates the image.
- ◆ The popular connection used by scanners and digital cameras is USB.
- ◆ Digital cameras are measured by the quality of image they provide. The measurement unit is megapixels — millions of pixels. The larger the number of megapixels, the higher the resolution and the larger the image.
- ◆ CODECs compress videos so that they don't use up as much hard disk space. In order to play a video on the computer, you need the same CODEC installed on your system that was used to compress the video.
- ◆ If you're having trouble with your microphone picking up the sound around you, make sure that the microphone is set as a recording device.

Prep Test

- 1** What type of scanner is portable and is swiped over the object to be scanned?
 - A Flatbed
 - B Handheld
 - C Sheetfeed
 - D Drum

- 2** What technology do flatbed scanners and digital cameras use to create an image?
 - A OCC
 - B CCO
 - C PMT
 - D CCD

- 3** What zoom type on a digital camera simply magnifies the image to emulate the zooming feature?
 - A Fixed zoom
 - B Digital zoom
 - C Optical zoom
 - D Pixel zoom

- 4** Digital video cameras typically use which type of connection to connect to the computer and upload the video?
 - A Serial
 - B Parallel
 - C FireWire
 - D USB 2.0

- 5** The standard used by musical equipment to communicate with one another or with a computer is known as what?
 - A OCC
 - B MIDI
 - C WAV
 - D Codec

6 Video files are compressed by using what?

- A Encryption key
- B MIDI
- C CCD
- D CODEC

7 What type of scanner uses a glass tube that emits light?

- A Flatbed
- B Handheld
- C Sheetfeed
- D Drum

Answers

- 1 B. A handheld scanner is carried around by the user, and the scanner is moved over the object to scan the object. *See "Understanding Scanners."*
- 2 D. *Charge-coupled device* (CCD) is the technology used by flatbed scanners and digital cameras to create the image. *Review "Understanding Scanners."*
- 3 B. Digital zooming is the zoom type that simply magnifies the image, just like you would do yourself in image-editing software. This could make the image appear grainy. *Check out "Digital cameras."*
- 4 C. Digital video cameras typically use IEEE 1394 (FireWire) connections to connect to the computer. *Peruse "Digital video cameras."*
- 5 B. *Musical Instrument Digital Interface* (MIDI) is the standard that allows musical instruments to communicate with one another or with the computer. *Take a look at "MIDI."*
- 6 D. A CODEC is used to compress and decompress video files. *Peek at "Understanding CODECs."*
- 7 D. A drum scanner uses a glass tube that emits light. *Look over "Understanding Scanners."*

Chapter 7: Dealing with Portable Computers

Exam Objectives

- ✓ Identifying portable computer components
- ✓ Understanding AC adapters
- ✓ Finding out about LCD panels
- ✓ Upgrading and expanding a laptop

In this chapter, you find out about portable computers and their supporting technologies. On the A+ exam, you're required to identify characteristics of a laptop, such as the types of batteries and displays, along with how to upgrade a laptop's hard drive or memory.

Laptops gained popularity years ago due to their small size and mobility. A standard desktop computer *can* be carried, but it's inconvenient and causes wear and tear on your biceps. Assume, though, that you are actually happy with the idea of carrying around a big, heavy desktop computer; the other problem with desktop computers is that you won't necessarily always have a power outlet to plug it into. A big benefit of the laptop is that it has a battery to use as a power source, which is great for when you're on a plane or out on safari and don't have a wall outlet.

Although laptops are the portable computer type to focus on for the A+ Certification exam, you should also be familiar with handheld devices such as *Personal Digital Assistants (PDAs)*, which a lot of people use to keep track of their daily schedule and contacts.

This chapter helps you identify, install, configure, and upgrade special portable computer components, such as batteries, hard drives, and memory. Procedures in this chapter guide you through the installation and removal of these components on laptop computers.

Identifying Portable Computer Components

Although you'll find a number of different laptop brands on the market, each make and model of laptop has common components. This section introduces you to some of the common components found in laptop computers. Figure 7-1 displays these components.




Figure 7-1:
Identifying
the major
components
of a portable
computer.

The following is a quick description of some of the major components found on laptop computers. Some of these components are also found on desktop computers and serve the same purpose for the laptop. For example, a video board built into a laptop does the same thing as a video board built into a desktop computer.

- ◆ **Battery:** The battery supplies power to the laptop and its components. The battery is charged or maintains its charge from the AC adapter that plugs the laptop into the wall outlet.
- ◆ **AC adapters:** The AC adapter supplies power to the battery so that the battery can either recharge or pass that power on to the laptop components.
- ◆ **DC controllers:** The DC controller is responsible for protecting devices such as cards, chips, adapters, and circuitry from power surges.
- ◆ **PCMCIA cards:** *Personal Computer Memory Card Industry Association* (PCMCIA) cards are used to add external drives, modems, and LAN adapters to laptop computers that didn't ship with one of these components. The PCMCIA cards expand on the computer's capabilities.
- ◆ **Video adapters:** The video board is responsible for converting the digital data from the system into information that the display can use to create the image.

- ◆ **LCD panels:** The primary display type on laptop computers is the liquid crystal display. The LCD panel is increasingly replacing CRT (cathode ray tube) monitors for desktop computers as well.
- ◆ **Keyboard:** The keyboard found on laptops works the same way that a keyboard works with desktop computers only the laptop keyboard is a smaller version of the keyboard found on desktop computers.
- ◆ **Hard drives:** Laptops ship with hard drives that are physically a lot smaller than the ones found in desktop computers and typically use less power and store less information.
- ◆ **Memory:** RAM is installed in laptop computers in a similar fashion to desktop computers. The only difference is the type of memory modules, known as *Small Outline Dual Inline Memory Modules (SODIMMs)*, used on laptop computers today.

Looking at Laptop Batteries

Laptop computers ship with a battery that can be removed or replaced at any time. The battery acts as the primary source of power for the laptop and all of its components. The battery maintains its charge, or is recharged, by the AC adapter that plugs the laptop into the wall. The following sections outline everything you need to know about portable-computer batteries for the A+ exam.

Different types of batteries

It is important to understand that a laptop computer has two types of batteries, a main battery and a CMOS battery:

- ◆ **Main battery:** The main battery is the battery that supplies power to the laptop and its components. A laptop is designed to work with one of four types of main batteries:
 - *Alkaline:* An alkaline battery is usually found in palmtop computers and is the same battery type you find in calculators.
 - *NiCad (Nickel-Cadmium):* A NiCad battery is typically found in laptop computers and is very heavy. The bonus of NiCad batteries is that they are fairly inexpensive. NiCad batteries typically need to be recharged after three or four hours of use.
 - *Li-Ion (Lithium-Ion):* A Li-Ion battery is a light battery that has a lot more battery charge time than NiCad batteries. A Li-Ion battery is also more expensive than a NiCad battery.
 - *NiMH (Nickel-Metal Hydride):* Nickel-Metal Hydride batteries are environmentally friendly because they do not contain toxic materials.

They are the same weight as NiCad batteries but are more expensive and don't last as long as a Li-Ion battery.

- ◆ **CMOS battery:** Just like a desktop computer, a laptop computer has a CMOS battery that is responsible for holding a charge to CMOS RAM so the CMOS configuration can be maintained.



On the exam, if you are asked which battery is the best battery for laptop computers, the answer is a Li-Ion battery because it provides the most durability and performance. It is more expensive than the other types of batteries, though.

Another type of battery that you may hear of from time to time is a *smart battery*. A smart battery has its own power circuitry, which is responsible for monitoring the battery performance, output voltage, and temperature, and also communicates the status of the battery back to the system. A smart battery is 15 percent more efficient than the other battery types and is also the most expensive type.



Handling batteries

Before you handle any batteries, be forewarned that batteries (except NiMH batteries) contain toxic chemicals that can cause harm if the battery explodes. Be sure to follow these guidelines when handling batteries:

- ◆ **Keep batteries away from fire and water.**
- ◆ **Never open or dismantle a battery.**
- ◆ **Try not to drop or throw a battery.**
- ◆ **Be sure to follow the manufacturer's guidelines when working with the battery.** For example, if you're storing the battery, check the documentation on suitable areas for storage.

Maximizing battery performance

One of my major pet peeves with batteries is that over time, the charge time of the battery fades away. When I purchased my new laptop, I could get about 3 hours out of the battery; now I'm looking at about an hour and a half. If you want to get the best performance out of your battery, follow the practices I outline below:

- ◆ **Be sure to fully charge and discharge a new battery or a battery that has not been in use for a while.** Completely discharge a battery to allow it to charge to its full potential. To discharge a battery, power on the portable computer without having it plugged in and leave it on overnight (or longer, if necessary) until the battery power is 100% used. You can then recharge the battery by using the AC adapter.

- ◆ **Regularly charge and discharge batteries completely every two to three weeks to keep them healthy.** Be aware that you need not do this for a Li-Ion battery.
- ◆ **Keep batteries clean.** This helps maintain a good connection between the battery and the portable device.
- ◆ **Don't leave batteries dormant for long periods of time.**
- ◆ **Store batteries well.** Be sure to store batteries in a cool, dry, clean place away from dangerous elements such as heat and other metallic objects.

Changing batteries

A number of laptop users find it useful to purchase an additional battery that they can keep charged and then switch with an existing battery that is losing its juice from use. I don't know how many times I've been on a 5-hour flight, and an hour and a half into the flight, my battery loses its charge. In this example, it would be useful to have another battery that I could switch over to.

To change a battery on a laptop, power the laptop down and turn the laptop over. You will notice a spot on the bottom of the laptop where the battery is placed. It will most likely have a battery symbol beside it with a lever that you slide to release the battery from its slot. Remove the old battery and then lightly place the new battery in the slot and clamp it in place by pressing down (as shown in Figure 7-2).




Figure 7-2:
Replacing
a battery
on a laptop
system.



Before you put in a new battery, make sure that you're using the right type of battery for the laptop. Checking the documentation for the system is the best way to do this.

Power management features

Laptop computers include *power management functions* that make the laptop “go to sleep” when there has been no input from the keyboard or mouse for a period of time. When the computer is in sleep mode, it still runs, but power-drawing features like the screen and hard drive are suspended until you “wake” your laptop by pressing a key. Powering down these processes when you're not using them helps preserve the life of the battery.

Two power-management standards have come out over the years — *Advanced Power Management (APM)* and *Advanced Configuration and Power Interface (ACPI)*. These power-management standards have a set of features that allows the system to conserve power, again resulting in less battery usage for laptop systems.

Both APM and ACPI are designed to accomplish the same goal, but ACPI is the newer power management standard of the two and is what most systems support today. APM was implemented as an *application programming interface (API)* in Windows, which allowed developers to call a library of code that controlled the power management features of the laptop. ACPI was implemented a little differently; it is a set of BIOS routines that control the power management features, which means that to configure ACPI, you make changes through CMOS.

Understanding AC Adapters

The *AC (Alternating Current) adapter* is nothing new to anyone who has worked with electronic devices. The AC adapter on a laptop, just like on an electric razor or a Game Boy, supplies power to the battery so that it can supply power to the laptop and all of its components. The AC adapter also recharges the battery so that you can use it when you are on the road. For example, I am writing this text in a coffee shop and running my laptop off my battery, which was charged up last night by the AC adapter. Figure 7-3 shows an AC adapter.



Be sure to always use the AC adapter that came with your laptop or one that is approved by the laptop manufacturer. Although other adapters may fit your computer, the voltage and amperage may not match the laptop and could cause damage.

Figure 7-3: Looking at an AC adapter that supplies power to the laptop and charges the battery.



AC adapter problems

When it comes to AC adapters, there are not a lot of issues that can arise. Two major problems that you will hit with AC adapters are listed below:

- ◆ **Cable problems:** One of the major issues that can arise when you travel a lot with a laptop is a break in the cable due to wrapping the cables up a lot. When you have cable problems, it's best to simply order a replacement adapter.
- ◆ **AC port damage:** Another major problem is that the AC port on the back of the laptop could get damaged. AC ports get damaged because the laptop is not being used on a hard, flat surface. A number of people I have talked to who have problems with the AC port not holding a connection to the AC adapter have been using their laptops to watch movies in bed with the laptops on their knees, which bang into the AC connector. In this scenario, it is best to have the laptop serviced by the manufacturer to ensure the correct electrical parts are replaced.

AC adapter troubleshooting and repair

If you find that you have a broken wire or faulty electrical components, then you can attempt to repair the problem yourself — but unless you have an electronics background, you should probably leave that to the experts and simply replace the AC adapter.

Learning about LCD Panels

A *Liquid Crystal Display (LCD)* is a popular display type that has been found in laptops for many years and is now a popular choice for display with

desktop computers and TVs. Liquid crystal displays also are popular on watches and other electrical devices.

An LCD has two sheets of material surrounding a liquid that contains crystals that act as pixels for the display. Each crystal has a red, green, and blue cell that is illuminated by an electrical charge hitting the crystal — which then creates the image we see on the screen.

There are two major types of LCDs: monochrome and color. There are also two subtypes of color LCD displays: active matrix and passive matrix:

- ◆ **Monochrome:** The original LCDs were single-color displays that showed blue or dark gray on a light gray background.
- ◆ **Color:** There are two forms of color display, active matrix and passive matrix:
 - *Active matrix:* An active matrix display, the most popular type of display today, uses at least one transistor per pixel, or crystal, which allows the electrical charge to be held longer on the crystal. This helps to create very crisp images with high resolution. With active matrix displays, the images are clear and easy to view, even from an angle. Due to the transistors, the active matrix display uses more power than a passive matrix display. Active matrix displays are also known as *Thin Film Transistor (TFT)* displays.
 - *Passive matrix:* A passive matrix display has one transistor for each vertical column of the display and one for each horizontal row of the display. The electrical signal is sent down the appropriate horizontal and vertical row of the display, and where they intersect is the pixel that is illuminated. Passive matrix doesn't have a transistor per pixel, so it doesn't use as much power as an active matrix, but it's slower to produce images and produces lower-quality images. Another popular term used to describe passive matrix is *dual-scan*.

Handling LCD panels

Because the LCD is so easily damaged, you should take care when handling it. The following are a number of key points you should follow to keep your LCD panel working well:

- ◆ **Make sure that you keep the temperature and humidity at acceptable levels.** The humidity should be below 60 percent, and the temperature should be no more than 40 degrees Celsius (that's 104 degrees Fahrenheit). This means that you should avoid leaving your laptop sealed up in your car on a hot summer day.

- ◆ **Use only a manufacturer-approved solution to clean your LCD.** Consult your owner's manual for a description of what fluid you should use to clean the display.
- ◆ **When cleaning the screen, don't scrub the surface vigorously or touch it with anything that has sharp edges.** To clean the display surface, be sure to wipe the screen gently with a soft cloth and manufacturer-approved cleaning fluid.

Connecting an LCD panel to a computer

The LCD panel is connected to your laptop by small hinges. To replace the old LCD panel with a new one, follow these steps:

1. **Although you should check the owner's manual or contact the manufacturer to find out exactly how the LCD panel is replaced on your laptop, in most laptops, you press the spring of the hinges gently to detach the old LCD panel.**
2. **After you have removed the old LCD panel, slide the new LCD panel firmly in place.**
3. **When you have the new LCD panel in place, be sure to close the panel and reopen it to verify that you fixed the LCD panel correctly. If it closes and opens easily, you replaced the LCD panel successfully.**



Because the LCD panel is the weakest part of the laptop computer, you want to be careful when handling the LCD and even opening and closing the laptop. If you damage the LCD, you need to replace it.

Understanding Laptop Input Devices

After you're familiar with the laptop's major components, such as batteries and displays, take a look at the myriad input devices used to get data into the laptop. The major input devices are the keyboard, touch pad, and rubber mouse ball located on the keyboard of some laptops.

Laptop keyboard

The *keyboard* is the most common input device used on both desktop systems and laptops. The keyboard allows you to communicate with the system by converting the keystrokes into corresponding letters and numbers. For information on the different types of keyboards and how they work, refer to Book III, Chapter 2.



For the exam, be aware that the Enhanced 101-key keyboard is the popular keyboard found on desktop and laptop computers today.

Maintaining keyboards

Over time, the keyboard collects a lot of dust and dirt behind the keys. It is important to clean keyboards periodically, and the following outlines the basic steps for cleaning a keyboard:

- 1. Hold the laptop at an angle and use compressed air to blow the dust and dirt from behind the keys.**
- 2. Clean the tops of the keys with a soft cloth dipped in a manufacturer-approved cleaner. You could also buy some wet cloths with the appropriate cleaning fluid already applied.**
- 3. Using a lint-free swab, clean any remaining dirt from between the keys.**

Handling keyboard problems

Like most replicable components today, replacing a keyboard is far more efficient than repairing it. The following steps provide a basic guideline you can use to troubleshoot any keyboard-related problems:

- 1.** If you are having troubles with the keyboard after replacing it, be sure that the keyboard is seated correctly.
- 2.** Reboot the computer. Yes, rebooting does work from time to time! Rebooting the computer reloads the device drivers, which are responsible for controlling the corresponding devices — in this case, the keyboard.
- 3.** If the keyboard still doesn't work, or if just some keys don't work, then you may need to replace the entire keyboard. Personally, I find that the keyboard is the first item to go on laptops.

Touch pad and the rubber mouse ball

Two other major input devices are the touch pad and the mouse ball, which are both responsible for controlling the mouse pointer on the laptop. Most laptop users use a normal mouse with their laptop by connecting the mouse to either the PS/2 port or a USB port.

If you don't want to use a PS/2 or USB mouse with your laptop, you will typically control the mouse pointer by using a touch pad or a mouse ball.

Touch pad

The touch pad, shown in Figure 7-4, is a very common component found on laptops today. It is a small, rectangular surface located in front of the keyboard. In order to use the touch pad to control the mouse pointer, you drag your finger over the surface, and the mouse pointer moves in the same direction.

Figure 7-4:
Using a
touch pad
to move the
mouse
pointer.



Below the touch pad are two buttons. These buttons act as the left and right mouse buttons. With most laptops, you can tap the touch pad instead of clicking the left button to perform a “click” operation, and you can tap the touch pad twice to double-click. If you find that the touch pad does not react to the tapping, have a look in the driver CD that came with the laptop to see if there is a driver that needs to be loaded for the touch pad.

Most laptops also have a button that can be used to turn the touch pad on and off. This can be useful if you’re going to use a USB mouse and don’t want the touch pad to interfere when you accidentally rub over it — which I find I do with my thumb as I type.

Rubber mouse ball

Instead of a touch pad, some older laptops shipped with a rubber mouse ball located in the middle of the keyboard. This rubber mouse ball was like a little joystick that controlled where the mouse pointer would go in the Desktop. Along with the rubber mouse ball, you would also have two buttons below the keyboard that would perform the left- and right-click mouse actions. Today’s laptops usually ship with the touch pad.

Laptop Communication Components

Laptops communicate with each other and with networks in a few different ways. In the following sections, I discuss some of the popular types of communication components found on a laptop.

Network card and modem

The most obvious communication component on the laptop is the built-in Ethernet network card. You use this network card to plug the laptop into a network via a typical UTP cable. Most laptops today ship with 10/100 Mbps network cards that allow for fast data transfer.

You can identify the integrated network card by looking on the side of the laptop for the RJ-45 jack, which looks similar to a telephone jack, only a tiny bit bigger.

Speaking of telephone jacks, the network port is typically located right beside a telephone port for the built-in modem on the laptop. Most laptops also indicate these ports with little pictures so that you know which port is the modem port and which is the network port.

Wireless network card

Most new laptops ship with a wireless network card as well as a wired network port. The wireless network card built into your laptop allows you to connect to a *wireless access point* and access the network and Internet without using a physical network cable.

Laptops today typically allow you to enable or disable the wireless network card quickly by pressing a button on the laptop. The wireless button on my laptop is in the top-right corner, but I have also seen the button on the front panel near the hard drive and power lights. From a troubleshooting point of view, you want to make sure that the wireless network card is enabled if you are trying to connect to a wireless network.

Other communications ports

A few other communication devices are popular in laptops today. They are Bluetooth, infrared, and WAN-cellular technologies. The following sections give a brief overview of each of these technologies. For more on these technologies, see Book VIII, Chapter 2.

Infrared

Infrared wireless communication uses infrared light signals to send data from one device to another. Infrared is popular with TV and VCR remote

controls, but as most of us have found out when using the remotes, infrared requires line of sight. *Line of sight* means that there can be nothing blocking the pathway of the light beam between the two devices. If the pathway is blocked, communication stops. Infrared communication is also limited by its bandwidth; it can deliver information at up to 4 Mbps.

Bluetooth

Bluetooth is a newer form of wireless communication that is often used to connect a portable computer to peripherals, such as a printer on the network. Bluetooth doesn't rely on line of sight because it doesn't send the data as light beams — it sends the data over a radio frequency and can be used only when the two devices are close together. Bluetooth has a maximum distance of 10 meters.



For the exam, remember that Bluetooth is popular in portable devices, such as handheld devices, and sends data at approximately 1 Mbps.

WAN-cellular

WAN-cellular is a technology that allows you to use your cell phone to make a network connection from your portable device to a remote location. The benefit of WAN-cellular technology is that you can connect your portable device to your network from anywhere that you can get a cellular connection. In order to use WAN-cellular technology, you need a digital cell phone that can connect to your portable devices, such as your handheld device.

Port Replicators and Docking Stations

If you use a laptop at the office a lot, you might want to consider using a port replicator or a docking station to get the functionality of a desktop PC from your laptop. A *port replicator* gives a laptop access to standard computer ports by connecting to an expansion port typically found on the back of the laptop. Once you connect the port replicator to the back of the laptop it provides serial, parallel, keyboard, and mouse connectors — allowing you to use those types of devices on the laptop even if the laptop itself does not contain the ports.

A number of companies set up docking stations for their laptop users. A *docking station* acts as a port replicator by having ports exposed out the back of the docking station. The docking station is different than the port replicator but the docking station also can contain features such as drive bays, network port, and desktop computer-sized expansion slots.

Expanding on a Laptop

Expanding on a laptop's capabilities usually involves upgrading the system memory or hard disk, which I discuss in the "Upgrading Your Laptop" section a bit later, or expanding by adding a *Personal Computer Memory Card Industry Association (PCMCIA)* card to the system. A PCMCIA card, also known as a *PC card*, is a credit-card-sized device that is inserted into a PCMCIA slot, usually on the side of the laptop. It allows you to expand on the laptop's capabilities. For example, you could add a network card, modem, memory, hard disk space, or maybe even a wireless network card to the laptop by adding a PCMCIA card (shown in Figure 7-5).



Figure 7-5:
Inserting a
PCMCIA
card into the
PCMCIA
slot.

Originally, PCMCIA cards were used only to add memory to the laptop, but today, you normally add memory by inserting it into a slot found on the bottom of the laptop — so the PCMCIA slots are now used for modems, network cards, and wireless network cards.



For the A+ Certification exam, be familiar with the three different types of PCMCIA cards:



- ◆ **Type I card:** A Type I card is 3.3mm thick, with a single row of sockets that connect into the slot. Type I cards were used to add memory to the laptop and are the original reason for the expansion architecture.
- ◆ **Type II card:** A Type II card is 5.5mm thick, with two rows of sockets that connect into the slot. Type II cards are usually network cards and modems that are added to a laptop.
- ◆ **Type III card:** A Type III card is 10.5mm thick, with three rows of sockets that connect into the slot. Type III cards are used to add hard disk drives.

Be sure you're familiar with the thickness of each card type and what types of devices are typically delivered in what card type.

Each of the types of cards has a corresponding PCMCIA socket and you can mix and match the cards into different sockets as long as the card fits in the socket. For example, a Type I card will go into a Type II socket but you will be unable to fit a Type III card into a Type II socket.

In order for the card to work in your laptop, the laptop has two services used to manage PCMCIA cards. The two services are as follows:

- ◆ **Socket services:** Socket services are BIOS-level routines that detect when a PCMCIA card is inserted into or removed from the system.
- ◆ **Card services:** Card services provide the interface between the card and the device driver. Card services also manage the assignment of system resources, such as I/O addresses and IRQs, to the PCMCIA card.

Adding a PC card

One of the most exciting features of PCMCIA is hot swapping. *Hot swapping* allows you to install a device while a computer is still running and have the device detected and configured by the operating system without rebooting. This is unlike adding a card to a desktop computer, where you typically have to shut the machine down, take the cover off, add the card, put the cover back on the computer, and finally start the computer up. With PCMCIA, you can insert or remove the card from the slot at any time.

Adding or installing a PCMCIA card is easy. Make sure you have the card oriented properly and then slide it into the slot until it fits firmly. If the card is installed correctly, the eject button beside the slot will pop out. When the card is inserted, Plug and Play should kick in and either load the driver or prompt you for the driver.

Removing a PC Card

To remove a PCMCIA card from a slot, you simply push the eject button, and the card will pop out. After you remove the card, make sure that the slot is

covered to keep dust from getting into the laptop. If the laptop doesn't have its own cover for the slot, you can purchase a *blanking plate*, which acts as an inserted dummy card.

Upgrading Your Laptop

In the following sections, I show you how to upgrade a laptop by adding components such as RAM and hard disk space. Makes and models of laptops vary, but the basic concepts of adding and removing components can be applied to most laptop systems today.

Upgrading memory

To add memory to the laptop, you simply flip the laptop over and find the compartment that holds the RAM; it is normally labeled with a symbol of a memory module, as opposed to the compartment that is labeled with what looks like a hard drive. If you are unsure what compartment to open you can check the documentation for the laptop.

Remember to ground yourself before touching any internal computer components! After you ground yourself, you can then remove the old SODIMM module and insert the new memory module. When inserting a SODIMM into a laptop, be sure to gently place the module in the socket at an angle (see Figure 7-6) and then clamp it down in place.

Figure 7-6: Installing a SODIMM into the laptop to add more memory to the laptop.



Adding hard disk space

Hard disk space is pretty scarce these days when storing media files like videos, MP3s, and pictures from digital camera. It doesn't take a lot of multimedia files (especially video) to fill up all of your laptop's memory.

You can add disk space to the laptop in a number of different ways:

- ◆ **Add an external USB drive:** One of the most popular methods of adding disk space to the laptop is to connect an external USB hard drive to the laptop. This lets you carry the data to a different PC as well.
- ◆ **Get a better hard drive:** On most laptops, you can actually replace the original hard drive with a manufacturer-approved hard drive that has more available space. For example, you could upgrade your 20GB hard drive to a 40GB hard drive. Most laptops have a panel on the back that can be removed to replace the hard drive (as shown in Figure 7-7).




Figure 7-7: Replacing the hard drive on a laptop.

- ◆ **Portable unique storage:** Portable unique storage devices can be carried around. These devices could be in the flavor of a USB external hard disk, or they could be a flash drive that plugs into a USB port (shown in Figure 7-8).

Figure 7-8: You can add storage space by using a USB external hard disk (top) or a USB flash drive (bottom).



Understanding Handheld Devices

Handheld devices are another type of portable device that is very popular today but does not fall into the laptop category. A handheld device, also known as a *Personal Digital Assistant (PDA)*, is much smaller than a laptop and uses touch-screen technology instead of a keyboard. You usually use a stylus to touch the screen and choose what you want to do with the handheld device.

Most handheld devices are used to keep track of personal information, such as your contacts or your schedule. PDAs today also have word processors and spreadsheet programs so that the business user can do some work on the PDA.

The PDA stores information in memory, or RAM, instead of on a disk like a laptop does. One of the benefits of storing the information in memory is that the PDA can be activated quickly with no boot-up time. Another benefit of a PDA is that the batteries typically last longer than the batteries found in laptop computers. A PDA battery can last anywhere from 12 hours to 2 months without being recharged, depending on the vendor. Figure 7-9 shows a PDA being used with the stylus.



Figure 7-9:
A PDA
being
controlled
by a stylus.

Most PDAs let you synchronize your calendar and contact list to a desktop PC by using a docking station. The PDA is placed into the docking station, and the content on the PDA is synchronized with the desktop computer.

There are many different makes of PDAs, but some of the popular ones are the Palm Pilot and Compaq's iPAQ. A number of vendors such as HP and Dell have their own versions of PDAs as well.

Getting an A+

In this chapter, I discuss the different components that make up a laptop. These components include items such as the battery, LCD, and keyboard. Be sure to remember the following facts about portable devices when preparing for the exam:

- ◆ The battery in a laptop is responsible for supplying power to the laptop and is charged by the AC adapter.
- ◆ There are two types of LCDs used in laptops: active matrix and passive matrix.

- ◆ Most laptops come with built-in ports, such as a network and modem port.
- ◆ To upgrade memory on the laptop, you can add a SODIMM to an empty memory socket located in the laptop.
- ◆ PCMCIA cards are used to expand on the laptop's capabilities.
 - *Type I cards* were used to add memory in the past.
 - *Type II cards* (the most popular these days) are used to add network cards and modems to the laptop.
 - *Type III cards* are used to add removable storage.

Prep Test

- 1** Which portable device is the most popular?
 - A Palm Pilot
 - B Laptop
 - C PDA
 - D Desktop computer

- 2** Which battery type is the best battery type for laptops?
 - A Alkaline
 - B NiCad
 - C Li-Ion
 - D NiMH

- 3** Which of the following components makes the laptop truly portable?
 - A Battery
 - B AC adapter
 - C DC controller
 - D PC card

- 4** Which of the following components supports hot swapping?
 - A Keyboard
 - B SODIMM
 - C PCMCIA cards
 - D None of the above

- 5** Which type of PCMCIA card is used by network cards and modems?
 - A Type I
 - B Type II
 - C Type III
 - D Type IV

- 6** Which of the following statements is true for socket services?
 - A Socket services is a set of BIOS routines that identifies when the card is inserted into the socket.
 - B It provides an interface between the card and the device driver when the card is inserted.
 - C It assigns I/O address and IRQ to the card.
 - D Both A and C.

7 What does LCD stand for?

- A Liquid Crystalline Display
- B Liquid Crystal Device
- C Logical Crystal Display
- D Liquid Crystal Display

8 Which statements are true for an active-matrix LCD? (Select all that apply.)

- A It is slow.
- B It displays dull images.
- C Each pixel is supported by its own transistor.
- D Images are easy to view from angles.
- E It displays low-quality pictures.

9 What type of memory module is used in laptop systems today?

- A SRAM
- B DIMM
- C SODIMM
- D SIMM

10 What does PDA stand for?

- A Personal Digital Architecture
- B Personal Digital Assistant
- C Performance Data Assistant
- D Performance Digital Architecture

Answers

- 1** B. Laptop computers are the most popular form of portable device. *See the chapter introduction.*
- 2** C. Li-Ion batteries are best suited for portable systems because they're lightweight and have a long lifetime. *Review "Different types of batteries."*
- 3** A. The battery is the component that most classifies the laptop as a portable device. *Check out "A Battery of Laptop Batteries."*
- 4** C. *Hot swapping* means that the component can be inserted and removed without shutting down the system; this is a characteristic of PCMCIA cards. *Peruse "Adding a PC card."*
- 5** B. Network cards and modems are Type II cards. *Take a look at "Expanding on a Laptop."*
- 6** A. Socket services are the BIOS routines that identify when a PCMCIA card is inserted or removed from the socket. *Peek at "Expanding on a Laptop."*
- 7** D. LCD stands for Liquid Crystal Display. *Look over "Learning about LCD Panels."*
- 8** C, D. Active-matrix LCDs produce high-quality images that are easy to read at angles. *Study "Learning about LCD Panels."*
- 9** C. Small Outline Dual Inline Memory Module (SODIMM) is the memory module type used in laptop systems today. *Refer to "Upgrading memory."*
- 10** B. PDA stands for *Personal Digital Assistant*. *Examine "Understanding Handheld Devices."*

Book IV

Maintenance and Troubleshooting

The 5th Wave By Rich Tennant



"Our automated response policy to a large company-wide data crash is to notify management, back up existing data and sell 90% of my shares in the company."

Contents at a Glance

Chapter 1: Performing Preventative Maintenance447

Chapter 2: Troubleshooting Procedures and Guidelines.....467

Chapter 1: Performing Preventative Maintenance

Exam Objectives

- ✓ Understanding the purpose and use of preventative maintenance tools and cleaners
- ✓ Identifying and applying maintenance techniques

When you buy a car you may take it in for regularly scheduled maintenance at regularly scheduled intervals, and if you do, you will have very few problems with your vehicle. On the other hand, if you only have repair work performed when there is a problem, you will find that after a few years, you always have a major component on your car having problems. Long-term, trouble-free use of any device involves taking preventative maintenance. Your computer is no different, and performing regular preventative maintenance will ensure trouble-free use. As a CompTIA A+ Certified Professional you need to know what preventative maintenance should be taken with your computer. This chapter takes a look at why you want to perform maintenance, what tools you can use, how to prevent problems by controlling the computer's environment, and maintenance of specific computer components.

Preventative Maintenance 101

Preventative maintenance is designed to prevent problems from occurring, unlike repairs, which is remedial or unscheduled maintenance. The main reasons to perform preventative maintenance include:

- ◆ **Save time.** Even though it takes time to perform some of the maintenance tasks, many can be automated. If a computer is well maintained, you will have fewer repairs to make, and the user will end up with less downtime.
- ◆ **Save money.** Time is money. Users who have less downtime related to IT activities will be more productive. Also, if you're not spending time repairing computers, then your time can be devoted to other pursuits that can improve the overall IT infrastructure.
- ◆ **Protect data.** Data is the life of many organizations. A company that has a major loss of computerized data has little better than a 40% chance of reopening, better than a 50% chance of going under within two years, and only a 6% chance of recovering and surviving long-term. With these

statistics in hand, you can easily see why many companies invest so heavily in backup technologies. But in many cases, companies don't plan backups, or they forget about the data that resides on desktop computers — even desktop and laptop computers should be protected.

- ◆ **Improve performance.** All computer systems seem to slow down over time and some simple steps can eliminate these slowdowns.



Most businesses can equate everything to saving money, but do not forget saving time, protecting data, and improving performance.

One of the biggest things that people don't do when it comes to preventative maintenance is create a schedule. When I bought my car, I got a schedule of maintenance activities and when they should be done.

Schedules are great, but if people don't follow them, what good are they? In addition to the schedule for my car, and so that I don't forget about it, I also get warning lights in the car that remind me to perform maintenance. In the case of the computer environment, the more processes that can be scheduled and automated the better, since if you rely on people to do the task, they will invariably forget to do the task occasionally, or regularly — so automate when possible.

The other important point is to not bite off more than you can chew — keep your maintenance manageable because if it is not, you probably won't do it. It's better to plan to do three things and get them done than to plan a dozen and get nothing done.

Your schedule for tasks will vary from mine, and may be affected by the industry you are in; but it is important to have a schedule for your tasks. If you do not have a schedule, your tasks will not get done. As I mention, schedules will vary from one organization to another, or even within an organization; for instance, a computer on a factory shop floor will likely need its internal components cleaned and its vents checked for obstructions more frequently than a climate-controlled office. A schedule will likely have daily, weekly, monthly, quarterly, and annually scheduled tasks. A list of possible tasks includes

- ◆ **Scanning for file system errors and viruses**
- ◆ **Updating virus and spam definitions and emergency boot disks**
- ◆ **Backing up data and CMOS settings**
- ◆ **Defragmenting hard drives**
- ◆ **Removing unnecessary hard drive files**
- ◆ **Cleaning screens, mice, keyboards, floppy and tape drives, and internal and external surfaces**

- ◆ **Ensuring that external system fans and vents, internal fans, and heat sinks are functioning properly and are free of dirt and debris**
- ◆ **Checking power protection devices and internal and external cable connections**



Although lists will vary, at a minimum you should remember these items on your list of preventative maintenance tasks.

As I mention, schedules vary depending on the computer's environment and use, but when looking at this list, you likely already have done some classification of daily, monthly, quarterly, and annual tasks.

With scheduled maintenance, you may have most tasks automated and performed remotely, but you should have an interval during which each computer is visited. Visiting the computer allows you to conduct visual and audio checks. This goes to the see no evil, hear no evil axiom. When you are near the computer, you may notice environmental issues such as piles of paper surrounding a computer and blocking intake and output ventilation holes, or you may hear a noisy fan that is close to failure — when fans are starting to fail, they are often unbalanced when they start up, and then they slowly level out.

Using Preventative Maintenance Tools

Your cache of preventative maintenance tools should include both software and hardware devices. A balanced approach to maintenance includes a good number from each category.

Materials and equipment

In order for you to be able to effectively perform preventative maintenance on your computer equipment, you will need to have the correct tools with you. These tools include the following:

- ◆ **Compressed air:** Used to remove debris and dust from surfaces that you don't want to touch or for dusting small corners.
- ◆ **Computer vacuum:** These vacuums can be used to remove dust and debris from a variety of surfaces within your computer, such as system fans.

A number of companies produce computer vacuums. One of the largest suppliers is Metropolitan Vacuum Cleaner, which makes the DataVac line (www.metrovacworld.com).

- ◆ **Cleaning cloths:** These should be lint-free cloths. If you use synthetic cloths, take care when you're working with cleaning chemicals; some chemicals can decompose the cloth or bleed out the dyes in them, which may leave residue on the surfaces you are trying to clean.

- ◆ **Cleaners and chemicals:** Most approved cleaners are mild, non-abrasive, and benzene- and ammonia-free. In most cases, chemicals should be avoided on most surfaces, and either a dry cloth or a cloth moistened with water should be used.
- ◆ **Cleaning pads:** These are soft, lint-free cloths that are already pre-moistened with mild cleaners. Cleaning pads are usually more of a convenience item than a necessity.
- ◆ **Manufacturer-approved cleaning kit:** Usually designed for things like floppy disk drives and tape drives.



Never spray cleaners near vents, around keys, or on most surfaces (best to just say never) of a computer and its peripherals due to liquid damage that can cause malfunctions or electrical arcing.

Software utilities

In addition to the hardware tools listed in the previous section, there are many software-based preventative maintenance tools, such as

- ◆ **Automatic update tools** for anti-virus, anti-spam, Windows, and so on.
- ◆ **Firmware updates**, although these should be applied only to respond to a problem because they sometimes create other problems.
- ◆ **Driver updates** can be thought of as firmware updates. If you don't have a reason to update, don't apply updates out of hand.
- ◆ **OS-level tools**, such as `defrag.exe`, `chkdsk.exe`, Task Scheduler (which are covered in Book 6, Chapter 4), and `ntbackup.exe` (which is covered in Book 7, Chapter 3).
- ◆ **Third-party maintenance tools**, which are often available as packages from large companies such as Symantec and McAfee or from smaller, independent companies.



Ensure that you remember hardware and software tools used for preventative maintenance.

Maintaining Environmental Controls

In addition to equipment-specific preventative steps, you can take several environmental steps to protect and preserve your computer. This section looks at those factors and elements.

Ventilation and airflow

When discussing equipment ventilation, only two factors come into play: external ventilation and internal ventilation. Poor ventilation can contribute

to higher-than-recommended CPU temperatures or other damage. Components that are affected by these internal and external factors include power supplies, processors, hard disk drives, and motherboard components like chipsets and capacitors.

External ventilation

The largest single variable in the area of external ventilation is actually the ambient room temperature. If the room is not cool enough, then your computer will never be cool enough, since the computer simply brings room air into the case to cool components, which is expected to be cooler than the air that is currently inside of the computer case. Room temperatures should be between 60 and 75 degrees Fahrenheit (15 to 25 degrees Celsius), but try to stay at the low end of that scale. Most computer manufacturers provide a recommended operating temperature range for their computers. If the computers are constantly used at higher temperatures, your failure rate will go up.

Just like extreme heat is bad, so is extreme cold, and many systems are exposed to extreme cold when being shipped to final destinations. If the outside temperature is substantially below your room temperature (especially if it is below freezing), then you should let your computer acclimatize to its surroundings before plugging it in and powering it up. The length of time depends on the temperature variation and whether this warm-up is done within or outside of the packing material. If the computer is left in the packaging materials, then the chance of condensation (which should be avoided) on internal components is reduced. In addition to condensation, thermal stress may damage the computer if it is heated too rapidly. This is similar to pouring boiling water into a cold glass — don't try this at home. Cold items tend to be more brittle than warm items.

Most computer cases have intake vents to bring cool air into the case, exhaust vents to expel the warm air, and are constructed out of the metal to dissipate additional heat that is not exhausted. If objects are placed around the case, they can block the vents and act as insulation for the case, preventing proper cooling of the computer from occurring. It's obvious enough that exhaust ports should remain open to allow heat to leave the inside of the computer, but if the intake ports are blocked, the computer does not get a supply of cool air and it will overheat. Without thinking, many people stack papers around and on their computer, blocking intakes.

People also push their computers back to free up desk space, thereby blocking rear exhaust ports. The computer placement and surrounding environment should be checked periodically to ensure that no airflow vents are blocked. In addition to this, many workstation areas built into desks do not take airflow into consideration, with solid doors on the front and only a small cable routing opening in the rear. Avoid placing your computer or monitor in any space that will reduce any sort of airflow to or from the unit.

Internal ventilation

Besides the environment around the computer, you can do a lot inside the case to aid or hinder ventilation or cooling. The rest of this section will look at factors that affect heat inside of the case, and what can be done to reduce heat or increase airflow. To start off, I should establish what the typical internal ventilation or cooling process is. As mentioned in the previous section, most cases are made of metal so that internal heat can be dissipated into the surrounding air. Also part of the case design (and some are better designed than others) are air intake areas and exhaust areas. For tower cases of all form factors, air usually comes in from the front, along the bottom, and is expelled at the top in the back, usually through the power supply.

Some power supplies don't have fans, so this exhaust is passive. But more often than not, the power supply has a fan that provides active exhaust. In either case, this situation takes hot air from the internal components and runs that hot air over the power supply components, adding to the thermal stress that the power supply's components already suffer from. To aid in the airflow, some power supplies have two fans, one on the external side and one on the internal side, to improve air flow, because if the airflow is increased, it is expected that each cubic foot of air that passes through the space will be cooler. The internal vents or fan on the power supply open either to the front of the case or to the bottom. When they open to the bottom, they are able to take more heat away from the CPU.

To keep the airflow going in the correct direction, keep the cover on and keep rear expansion slot covers in place. Also make sure that you keep all front ventilation openings clear of obstructions. If you don't do this, then the airflow won't go in the direction that the case designer planned.

Many case designers add extra fans and vents to aid in airflow. Standard vent locations include an exhaust location on the back of the case below the power supply, which can take a fan to exhaust rising hot air before the power supply has to deal with it. Intake fans are also popular in the lower front of the case, increasing the air intake function. Typical case airflow is illustrated in Figure 1-1, with a power supply having internal vents opening to the front of the case.

Because the fan on the heat sink typically blows air down onto the heat sink, helping with heat dissipation, another common area for vents and fans is on the side of case parallel with the heat sink fan, so that cooler air from outside of the case can be drawn directly to the CPU rather than moving the warmer air existing in the case over the CPU. To augment this airflow even further, some cases also include a funnel on the side of the case to directly channel the air to the heat sink fan.

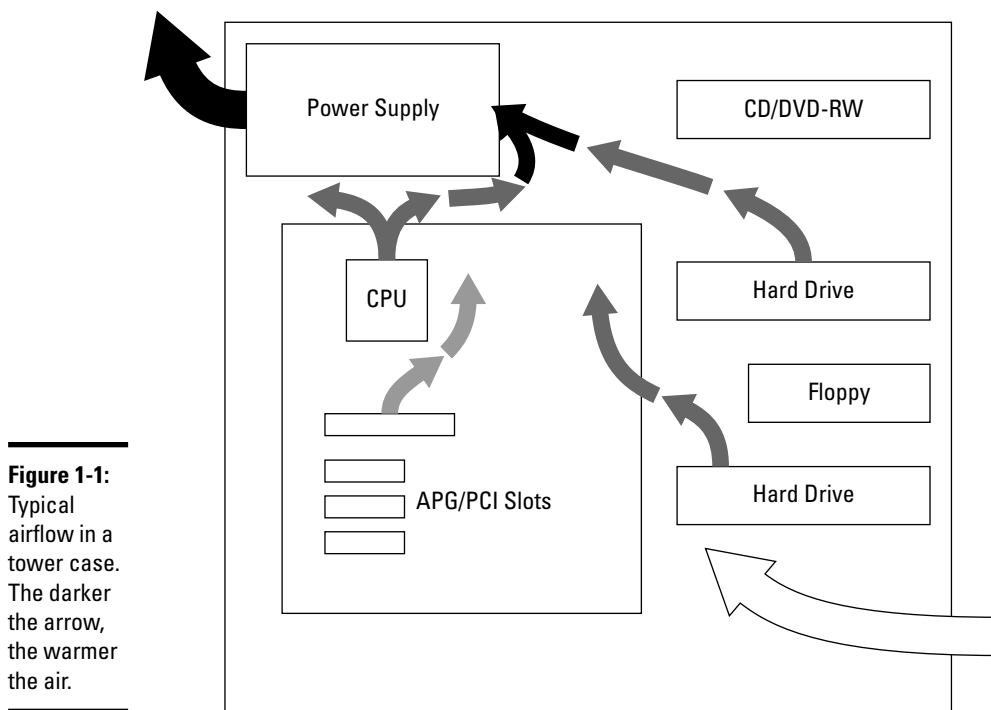


Figure 1-1:
 Typical airflow in a tower case. The darker the arrow, the warmer the air.

Less popular for most case designs, but sometimes added as an after-market addition, is the *blowhole*, which puts an exhaust port and a fan on the top of the case, drawing all the rising heat directly out of the case. If heat is a major problem, and this modification makes sense for you, it does give a lot of bang for its buck.

Internal airflow can also be improved by using a larger case, which just has more room for air to move around. There is no bigger heat challenge than taking a new multi-gigahertz processor, which generates more heat than slower processors, and adding it to a MicroATX case with 2GB of RAM and a couple of 7200 RPM IDE hard drives. This is a recipe for thermal challenge, especially if noise reduction is also a goal because noise reduction usually means reducing the number or the speed of the system fans. By taking the same components and putting them in a full tower case, the additional air flow from the larger case will reduce the overall system heat.



Reducing the number of ribbon cables used on IDE devices can also improve airflow. These ribbon cables make little walls within the case that air has to go around. Ribbon cables can be replaced by rounded cables, or the drives can be replaced with SATA drives, which have smaller cables to begin with.

In addition the items already mentioned, wise placement of internal components can help as well. If you're adding two hard drives to the computer, then placing them in adjacent bays may seem to make everything look nice and neat, but it also reduces the airflow between the drives, and heat generated by the bottom drive will rise and increase the overall temperature in the top drive. Leaving the one-bay gap allows the system fans to draw some of that heat away from between the two drives, lowering the temperature of the top drive.

Humidity and liquids

Humidity is a double-edged sword and should be approached carefully. To reduce the chance of ESD, it is recommended to have humidity levels of 60 percent, but to reduce corrosion; it is recommended that humidity levels be as low as 35 percent. Areas that are too damp add to the corrosion of metal components on the motherboard, such as leads for capacitors and the hundreds of metal contact points in the computer. As if corrosion weren't bad enough, condensation is also an issue that can occur if that damp air is also warm as it makes contact with the cooler computer components. Because water is such a good conductor of electricity, condensation is never good and can lead to major electrical problems. Whenever possible, reduce the computer's exposure to overly humid areas, which can sometimes be done by moving the computer to into the next room, and extending the cables for the monitor, keyboard, and mouse into the damp environment.

Just as overly damp environments are bad, extremely dry environments increase the risk of static build-up and cause early wear on capacitors, rubber rollers, and some other components.

In addition to damp environments, spills are another source of liquid problems. In general, avoid using liquids around the computer. If keeping liquids away from the computer isn't possible (who am I to deny you your morning caffeine?), then perhaps you can remove the computer from the area of the liquid. You can get special computer cases, keyboards, and accessories that have been hardened to protect them from moist environments, but often at that cost of cooling. To moisture-proof a case, most of the unit, including the ventilation openings, is covered with a rubberized material.

ESD

ElectroStatic Discharge (ESD) is caused when an electrical charge builds up on one surface, and then that surface comes in contact with another surface that has a lower charge. In most cases, your body builds up a charge — when any synthetic materials run against each other, such as your clothing, or moving across a carpet — and then when you come close to touching the computer, which is plugged into a grounded or bonded source, the electricity naturally attempts to balance itself out, often jumping the remaining gap and producing a surge of electricity and a spark. The discharge is normally

absorbed by the device, and the extra electrical power is drained through the grounding device, but if the gap was formed as you were inserting a piece of RAM or an expansion card, then you could damage the slot or the chipset it is connected to.

Because static is more prevalent in dry locations, you can increase the relative humidity in the area to reduce the chance of static discharge. When servicing a computer, use anti-static work mats and straps to keep you and the device you are working on at the same potential energy. Even though damage rarely results from small ESD, there is no sense in taking needless chances.



Do not forget the relationship between ESD and humidity. Lower ESD risks by increasing humidity.

Dirt and dust

Users are often surprised by how much dust and dirt collect inside a computer even though the room seems dust-free. Computers collect even more dust when they're placed on the floor, where they essentially act as air filters. In most cases, the computer should be cleaned out annually, but if it has to be placed in a dirty or dusty location, such as a factory floor, then internal cleaning should be scheduled more frequently. Just like water-hardened cases can be purchased, special cases are available for dirty and dusty areas, also again usually at the cost of system cooling.

When the system components are clear of dust and grime, they last longer and function properly. On several occasions, I have attempted to install a Windows OS only to get errors during the file copy phase. Normally, this would suggest bad media, but it has invariably been caused by a thick layer of dust and dirt very visibly covering all components. After I use compressed air to get rid of the grunge, the file copy problem invariably goes away. In addition to the dust causing various short circuits, it also causes heat to be retained by the components that are covered in the dust and dirt.

EMI

ElectroMagnetic Interference (EMI), also known as *Radio Frequency Interference (RFI)*, is caused by rapidly changing signals in electrical circuits radiating a field around the circuit. This may be caused naturally as part of the circuit's normal operation, or it may indicate a problem with the installation or operation of the circuit. This interference often affects AM frequency bands and other low-frequency devices, such as speakers. If you have been in a boardroom with a speaker phone or radio nearby, you will often hear EMI caused by ring signals being sent to the occupants' cell phones. These signals precede the phone ringing and usually cause a static noise on the speaker or radio. EMI caused by cell phones is the reason why cell phones are banned in certain areas of hospitals — they can disrupt electronic medical devices within a meter or yard.

Solutions to EMI include physical isolation of devices, installation of dedicated circuits, and power conditioning of incoming power.

Power, UPS, and suppressors

Power conditioning is cleaning or regulating power that is supplied to devices. In addition to conditioning power and reducing EMI, surge suppressors and UPS can provide protection to your computers and electronic devices. If you are using these power conditioning devices, it is important to check the status of these devices periodically. Most major UPS have self-test buttons or automatic tests that can be scheduled and reported to administrators, as well as reporting on load levels and battery health. To help UPS and other power protection devices, the use of dedicated circuits reduces the chance of other devices on the circuit causing EMI, other line noise, power surges, or circuit overloads.

In addition to keeping the power clean and available, grounding of equipment is important. All computer equipment comes with polarity base plugs (with one side larger than the other), and most have ground connectors (the third connector on the plug). Some users who need extension cords will cheat and get two-conductor cords, thereby removing the grounding facility for their equipment. This should always be avoided due the risk of damage to equipment from static discharge.

If power interruptions or blackouts are expected, as happens in some areas of the country regularly, then computer equipment should be protected by a UPS or shut down when not in use. This reduces the chance of data corruption due to power loss. If electrical storms are a risk, then equipment should be protected by UPS or unplugged, preventing damage that might be caused by a lightning strike.

Aside from the previous points, many people ask, “Should I usually leave my computer on or shut it off?” There are two schools of thought on that, which should not be a surprise. Some factors that affect this decision are:

- ◆ **Convenience of having the device ready to work:** If you visit your computer regularly, then you may want to keep it on for this reason.
- ◆ **Power consumption:** Running your computer 24/7 will consume more power than shutting it down when you aren’t using it. Many people shut down systems that won’t be used again for 8 or 24 hours. This means either shutting down every night or shutting down just for weekends.
- ◆ **Thermal stress:** Thermal stress happens every time components heat or cool because they expand and contract. This expansion and contraction causes failures to occur. Leaving equipment on for longer periods reduces this problem.

- ◆ **Wear-out:** Most equipment has an expected lifespan measured in hours or days. This rating is the Mean Time Before Failure (MTBF). If components are forced to run continuously then they will hit their usage limits more quickly, so this problem is reduced by powering items down or having them enter a low power mode.

The real answer is that you should find the solution that works better for you and your organization.

Completing Maintenance Tasks

Now that you have taken a look at the basic principles of preventative maintenance, take a look at what can be done for specific components. This section will look at each of the major components that make a modern computer, and look at what tasks should be taken to properly maintain each of them.

Case and components

To keep most components in good health, regularly perform maintenance on your systems. The frequency for these tasks depends on the task, with external tasks being performed daily, weekly, or monthly, and internal tasks being performed monthly, quarterly, or annually. The following tasks should be part of your maintenance routine:

- ◆ Make sure you have good airflow both inside and around the case — this will help prolong the system life.
- ◆ Periodically clean the inside with a vacuum or compressed air to ensure that dust and dirt haven't built up on the components or fans.
- ◆ Verify that all the fans are functioning.
- ◆ Check the processor temperature by using internal sensor chips, relative temperature by touch, or a digital infrared thermometer.
- ◆ Ensure that the heat sink and fan are clean and free of dirt and debris.
- ◆ Check the position of expansion cards and removable chips to verify that they are correctly seated. Newer technology for slots has reduced the chance of chip creep and other issues.

Chip creep is caused by thermal expansion and contraction which happens during the heating and cooling of components, and causes components such as removable chips to shift position or work their way out of their mounting clips.

- ◆ Make sure external and internal cables are properly seated and connected. Some of these cables may work loose from movement of the case or thermal expansion.

- ◆ If you have systems that use custom CMOS settings, which is less of an issue these days, ensure that you have a good backup of the CMOS settings.

Power supplies

Very few maintenance tasks are associated specifically with power supplies:

- ◆ Verify that your power protection devices are working properly to ensure system integrity.
- ◆ Ensure that power outputs are correct. This can be verified with the use of a multi-meter or power supply testing tool. Both of these methods are covered in Book 4, Chapter 2.
- ◆ If system components have changed, verify that the power output is still sufficient.
- ◆ Make sure that the fans are clean.

Monitors

With their lack of moving parts, and low number of failures, monitors are often overlooked during the maintenance process. Even though they are often trouble free, you should include them in your maintenance routine, to ensure that they will provide trouble-free operation.

- ◆ Ensure that the monitor or OS power-saving feature is turning the monitor off properly, which not only saves power but also prevents image burn-in of a static image on the screen.
- ◆ Make sure the monitor has proper ventilation; ensure that the vents on the unit have not been covered. With traditional CRT (Cathode Ray Tube) monitors having large mostly flat tops, I have often seen users make it an extra filing area for papers, blocking the upper exhaust openings, or users with cluttered desks filling in the area around the base, blocking up the lower air intake openings. LCD (Liquid Crystal Display) monitors do not generate the same heat, but ventilation is important to get rid of the heat that they do generate.
- ◆ The screens of many monitors these days have special anti-glare coatings that are not supposed to be cleaned with standard glass-cleaning products. Always check the owner's manual for the correct cleaning solutions for your monitor. When in doubt, use a dry cloth or a cloth moistened with water for troublesome spots.



Always apply cleaners to a cloth, not directly to the monitor. This warning should also be applied to any computer equipment that you are cleaning with liquid cleaners.

Keyboards and mice

For keyboards and mice, preventative maintenance includes checking that cables are not damaged by keyboard trays or other external forces. Verify that they are properly connected and have not been pulled loose.

For keyboards, proper maintenance consists of the following:

- ◆ Ensure that food and liquids are kept away from the keyboard's spill area, which is usually about twelve inches or thirty centimeters.
- ◆ If a key's surface appears dirty, you can clean key caps with a cloth moistened with a mild cleaner or water. You should hold the keyboard upside down during this process to keep any debris from falling between the keys.
- ◆ If debris between the keys is causing key motion problems, then some compressed air can blow it out.
- ◆ If there is no issue, you may want to leave the status quo. Using compressed air may move debris to a position that causes an issue.



How you keep your mice clean and happy depends on what type of mice you have:

- ◆ For optical mice, verify that the optics are clean and that the surface that it is used on is not too glossy. Some tabletops and mouse pads may not be suitable for optical mice, so add or remove a non-glossy mouse pad as needed for mouse problems.
- ◆ For mechanical mice,
 - Remove the retaining disc and the mouse ball and remove any dirt that has collected on the rollers inside the mouse by using a cotton swab moistened in alcohol.
 - Verify that the mouse pad is clean, and if it is cloth, replace it regularly because they tend to build up oils and other little nasties from your hand, and dust from the surroundings, which then build up on the mouse rollers.

Drives

Regular preventative maintenance for all types of drives includes verifying that cables are properly connected and sufficient power is being supplied.

Hard drives

For hard drives specifically, regular preventative maintenance includes

- ◆ Running Check Disk scans of your hard drive for errors
See Book VI, Chapter 4, to find out more about Check Disk (`chkdsk.exe`).

- ◆ Scanning for viruses with an anti-virus program, which, together with regular backups of important data, reduces the chance of problems that could affect the integrity of your data
- ◆ Defragmenting your hard drive to keep disk access at an optimal level
- ◆ Regularly monitoring the temperature and free space of your hard drives and CPU

You can monitor your hard drives locally with tools like SpeedFan (www.almico.com/speedfan.php), or remotely with SNMP (Simple Network Management Protocol — the TCP/IP device management and monitor protocol) tools like MRTG (www.mrtg.com) and Cacti (www.cacti.net). These tools gather temperature readings from your *S.M.A.R.T.* (*Self-Monitoring, Analysis and Reporting Technology*) hard drive, which is incorporated into most new drives. In addition, you can gather S.M.A.R.T. messages from your Event Viewer.

Locally, Windows XP provides warning messages when hard drive space is very low.

- ◆ Avoiding vibration and shock

Vibration and shock to the hard drive should be avoided because either may cause substantial data errors.

Floppy and tape drives

Floppy drives can be maintained with the following actions:

- ◆ Regular head cleaning with head cleaning kits based on manufacturer recommendations for frequency
- ◆ Verifying head alignment

To verify the drive's head alignment, format a new disk and then try to read it in other floppy drives. If the disk is unreadable, you might suspect an alignment problem in one of the drives.

When dealing with tape drives, you also want to clean the heads with the manufacturer-recommended cleaning tools, and this should be done based on the manufacturer's recommended frequency, which is usually based on hours of drive use. Verification on data recovery should also be regularly tested by attempting to read the data from the media.

CD drives

CD drives are optical drives, so occasionally the optical components inside the drive may get dirty, but this is rare when compared to media problems, which is usually where the problems arise. There are cleaning kits for the optical components, and you should check with your drive manufacturer for recommended use.

Storage media

To have consistent access to data you've recorded on storage media, you need to make sure that you protect that media:

- ◆ If your optical media is scratched or pitted to a point that it is unreadable, then it may be repaired or polished by using a commercial repair kit or Brasso metal polish.
- ◆ For very deep scratches in optical media, see if your local video store has a repair service.
- ◆ CDs can be damaged by glue found in many types of labels, so a felt-tipped marker is often the safest choice for labeling the CD.
- ◆ For all media, you should avoid physical contact with the media surface. Oil and dirt from your skin can damage the material immediately or over time.

Each type of media has a recommend shelf life, as estimated in Table 1-1; your mileage may vary based on media quality and storage conditions. All media should be kept in a clean, dry, cool, low-light place, and if magnetic media, away from electromagnetic fields.

Table 1-1 Shelf Life for Recorded and Stored Media		
<i>Media Type</i>	<i>Examples</i>	<i>Shelf Life</i>
Flexible Magnetic Media	Floppy disks, magnetic tape, Bernoulli cartridge drives	10 years
Rigid Magnetic Media	Hard disk drives, removable cartridge drives	15 years
Writable Optical Media	CD-RW, DVD-RW	5–10 years

Laptops

Laptops have their own maintenance issues, including screen types, battery usage, charging, storage, and theft.

Before cleaning a laptop screen, check the owner's manual for proper cleaning instructions. Many laptop LCD screens have special coatings or mesh filters that can be damaged if you don't know what you're doing.

With each new generation of laptop computers comes a new generation of batteries, each with its recommended care procedures. Nickel-cadmium (NiCad) batteries were replaced by nickel-metal hydride (NiMH) batteries, which were replaced by lithium-ion (Li-Ion) batteries, which were replaced by lithium-ion polymer (Li-Poly) batteries. The latest generations of batteries do not develop charge memories when not fully discharged between charges, cannot be overcharged, and have a better weight-to-power ratio.

One drawback of lithium-based batteries is that at room temperature, they lose 20 percent of their charge capacity per year, starting at the time of manufacturing; so a three-year-old lithium battery will have a maximum charge which is about 40 percent of what it was when the battery was purchased. This charge loss can be reduced to 6 percent per year if the battery is stored at 0 degrees Celsius or 32 degrees Fahrenheit — but then the battery needs to be warmed for over 24 hours to return to a usable temperature.

With Li-Ion and Li-Poly batteries, you no longer need to fully discharge batteries. In fact, it is best to keep them topped off, and deep discharges can reduce their life.

To keep your battery in tip-top shape, you should let the battery run down (except for Li-Ion and Li-Poly batteries) and perform a charging cycle every few weeks. In general, you can expect to replace your battery every four to five years. If your laptop is going to be used while connected to a power source, some manufacturers recommend removing the battery. See your owner's documentation to find out if this applies to you.

When moving your laptop between locations, you should use a travel case that is sized and designed appropriately for your computer. This case serves two main purposes: to keep the system free from dust and dirt and to protect it from shock.

Because laptops are so portable, they are ideal targets for theft, and all come with a lockdown point that you can use with a computer lockdown cable or motion sensor. The lockdown cable is a thick vinyl coated braided metal cable, like the cables used for bicycle locks, and needs to be attached to an immovable object, while the motion sensor sounds a loud alarm when the laptop is moved. You can see samples of each of these at the Targus Web site (www.targus.com/us/accessories_security.asp). One of these two types of devices should be used for a laptop when you are traveling, as well as in your office.

Printers and scanners

Dust, dirt, and debris tend to be the three big issues with printers. Much of this can be removed with a vacuum, compressed air, and a cleaning cloth, but here are some other cleanliness issues to take into account to keep your printers running smoothly:

- ◆ Printers have pickup rollers that can become dirty and should be cleaned with cotton swabs moistened with alcohol.
- ◆ The print heads on inkjet printers can become clogged and should be cleaned through built-in cleaning procedures on the printer (see your

owner's manual for the exact procedure for your printer), and if the printer has its print head as part of the ink cartridge, you can also soak the heads in warm water to clear clogs.

- ◆ Laser printers often have fans and vents that should be checked and cleaned out.
- ◆ Many laser printers have maintenance kits that should be purchased and installed as necessary. These kits often replace components that wear out, such as pickup rollers.
- ◆ Many printers also have a calibration routine and a test or diagnostics page that ensures that multi-color print heads are all in alignment and that all print mechanisms are functioning. As part of your maintenance routine, you should print the diagnostics page.

Scanners, like printers, can suffer from the dust, dirt, and debris issues. Scanners tend to have a sealed body, which keeps the internal optics clean, but they usually need to be opened every year or two to clean the internal components. During regular maintenance, only the deck (the glass *platen*) of the scanner needs to be cleaned; use a glass-grade cleaner unless your scanner has a special coating. In addition to this issue, scanners usually have a locking mechanism that should be used when moving the scanner, which will keep the scanner properly calibrated.

Because both of these devices suffer from a build-up of dust and dirt, you should try to place them in areas that are clean and dust free. I have seen people save money by not purchasing a printer stand, and leaving the printer on a carpeted floor; only to pay more than the cost of a stand in additional cleaning and repair services.

Getting an A+

In this chapter, you should remember

- ◆ Performing preventative maintenance saves time and money.
- ◆ Conduct preventative maintenance tasks at scheduled intervals.
- ◆ Standard tools for preventative maintenance, include cleaning materials, vacuums, compressed air, and software tools, such as `defrag.exe` and `chkdsk.exe`.
- ◆ Factors that affect computer lifespan are heat and humidity, ESD, dirt and dust, and incoming power issues.
- ◆ Each major component has specific maintenance that can or should be carried out.

Prep Test

- 1 Which of the following is *not* an acceptable cleaner for computer components?**
 - A Alcohol and cotton swabs
 - B Dry, lint-free cloth
 - C Ammonia-based cleaner
 - D Cloth moistened with water
- 2 Which of the following does *not* help improve computer airflow?**
 - A Adding exhaust fans
 - B Leave spaces between devices
 - C Keeping vents clear of debris and dust
 - D Using ribbon cables where possible
- 3 Which of the following is *not* a reason to implement a preventative maintenance schedule?**
 - A Reduce office temperature
 - B Improve performance
 - C Protect data
 - D Improve productivity
- 4 Which two items will help ensure that your preventative maintenance plans are followed?**
 - A Filing completed checklists
 - B Keeping the number of tasks limited
 - C Automating tasks
 - D Purchasing colored index tabs for files
- 5 What tools will help with hard drive maintenance? (Choose two.)**
 - A defrag.exe
 - B fdisk.exe
 - C chkdsk.exe
 - D mem.exe

- 6** What is one of the problems caused by low relative humidity?
- A EMI
 - B PKI
 - C ESD
 - D TCP
- 7** Which of the following can become an issue if you keep your computer running all the time?
- A Thermal expansion
 - B Wear-out
 - C Chip creep
 - D EMI
- 8** What factors affect the storage life of computer media? (Choose two.)
- A Media length
 - B Age of facility
 - C Temperature
 - D Humidity
- 9** What is usually found in vendor-supplied laser printer maintenance kits?
- A Paper pickup rollers
 - B Replacement power cable
 - C Paper
 - D Toner cartridges

Answers

- 1 C. Ammonia- and benzene-based cleaners are too strong for most computer components. While alcohol is not used as often as dry or water-moistened cloths, it is still acceptable for many computer-related components. *See "Materials and equipment."*
- 2 D. Ribbon cables create small walls in the path of airflow and can be replaced with rounded cables to improve airflow in most situations. *Review "Internal ventilation."*
- 3 A. Reducing office temperature is something that may help systems as part of a preventative maintenance plan, but it is not a reason to implement the plan. *Check out "Preventative Maintenance 101."*
- 4 B, C. Keeping the number of tasks limited keeps the process from becoming unwieldy. Automating tasks removes the need for action on a user's part, so that will also promote having the tasks performed. Filing and indexing physical files will not promote having the tasks completed but can be used to report on what has been completed. *Peruse "Preventative Maintenance 101."*
- 5 A, C. `defrag.exe` and `chkdsk.exe` are both tools that can be used to perform standard maintenance on your hard drives. *Take a look at "Software utilities."*
- 6 C. When relative humidity is too low, there is an increased risk of static build-up and electrostatic discharge (ESD). *Peek at "Humidity and liquids."*
- 7 B. When devices are running all the time, they can wear out. *Look over "Power, UPS, and suppressors."*
- 8 C, D. Media's storage life is greatly influenced by its storage environment, which ideally will be cool, dry, properly lit for archival purposes, and free of EMI sources. *Study "Storage media."*
- 9 A. Most laser printer maintenance kits include replacements for components that wear out, such as pickup rollers, as well as cleaning materials for the inside of the printer. *Refer to "Printers."*

Chapter 2: Troubleshooting Procedures and Guidelines

Exam Objectives

- ✓ **Identifying uses and purpose for standard troubleshooting tools**
- ✓ **Understanding and applying troubleshooting techniques to isolate and resolve problems for major computer system components**

In the preceding chapter, you take a look at preventative maintenance routines for your system's devices. But no matter how much preventative maintenance you perform, something will eventually go wrong. This chapter continues on with the topic of troubleshooting devices after things have gone wrong.

This chapter adds to the basic troubleshooting concepts that are introduced in Book I, Chapter 2. As a CompTIA A+ Certified Professional, you will be required to repair components from many areas of the computer system. This chapter covers the hardware and software tools of the trade that will help you complete the troubleshooting tasks quickly. You will also review the major components found in computer systems and address any troubleshooting steps related to those components.

Identifying Troubleshooting Tools

A good troubleshooting arsenal contains many weapons, both hardware and software. Not every computer problem is related to the hardware, and even when it is, software tools can sometimes help with the diagnosis. If you are doing field support, you should make a troubleshooting kit that has all of the tools that you use most often in a case that makes it easy to bring your tools to any location.

Hardware tools

In order to properly troubleshoot equipment-related issues, you will want to use the right tool for the job. This section will take a look at the hardware tools you should have, and what jobs you will perform with them.

Multimeter

You could buy a meter to measure voltage (volts), a meter to measure resistance (ohms), and a meter to measure continuity and current (milliamps and amps). Or you could just buy a *multimeter*, which is a combination of all of these different types of meters.



Make sure that you know what the multimeter is capable of measuring, since you will need to identify the proper times to use one, but not details of its operation.

You can find both digital and analog multimeters, and your choice is based solely on personal preference, but many people find that digital multimeters are easier to read. Also, when testing resistance in circuits, digital meters use only 1.5V rather than 9V, which is less likely to damage the circuit.

Both digital and analog multimeters are shown in Figure 2-1. The meter usually has a dial that lets you choose what you want to measure and the scale that you are measuring on. You should always use a scale higher than the reading that you are taking to avoid damaging the multimeter. The closer the scale is to the value you are attempting to measure, the more accurate your reading will be.




Figure 2-1: Analog and digital multimeters can both be used to test many system components.

Some other features to look for in a multimeter are:

- ◆ **Size:** Large meters are usually more feature-rich, but, for computer work, many of the smaller meters will be more convenient to transport with your troubleshooting gear.
- ◆ **Overload protection:** Protects you when voltage or current exceeds what the meter is capable of.
- ◆ **Auto ranging:** Automatically selects the appropriate scale for you but usually allows for manual override.
- ◆ **Detachable probe leads:** Allows for greater flexible by substituting different leads for different jobs.
- ◆ **Audible continuity test:** Allows you to perform the test without having to look at the meter.
- ◆ **Automatic power off:** Conserves battery life in the meter. It's a pain when you discover that the meter was left on and the battery has died.
- ◆ **Automatic display hold:** Keeps the latest stable reading visible on the display, making it easier to record the value if you are not able to see the meter when you are using the probes.
- ◆ **Minimum and maximum trap:** Records the highest and lowest readings in the same manner as the automatic display hold.



When measuring resistance, you should always remove the power from the device you are checking, unlike when you're checking voltage, which is done when the device has power.

In "Power supplies and batteries," later in this chapter, you see how to use a multimeter to verify voltages coming from power supply hard drive connectors.

Anti-static mat and strap

I discuss *ElectroStatic Discharge (ESD)* in the preceding chapter, and one of the best solutions to avoid ESD is to use a grounding or bonding strap, as shown in Figure 2-2. In addition to wrist straps, you can get anti-static mats for your workbench or floor. Both mats and wrist straps are made of non-conductive material and can be grounded to prevent charge build-up.



Taking ESD reduction steps, such as using a wrist strap, is always required when working on computer equipment.

Figure 2-2:
An anti-static wrist strap with a removable clip, allowing the plug to be used with some hardware.



Dealing with screws

This section will introduce you to the screws that are common to computer systems, as well as cover problems that you might run into with screws, such as dropping them.

Many types of pickup tools can be helpful when that screw, jumper, or other component slips from your fingertips. In some cases, you have no other means of retrieving that component from where it has fallen. There are two varieties of pickup tools — mechanical and magnetic — and both are inexpensive and have their uses. As long as you have one in your kit, you will be well off.



Some magnetic pickup tools, such as floor sweepers, have extremely high power magnets that could damage hard drives, but most small pickup tools use low power magnets that are safe near hard drives but that should not be placed next to floppy disk media.

To go along with the pickup tool for small bits and pieces, tweezers and needle-nosed pliers are beneficial when it comes to removing and placing small elements like jumpers.

Even though a screwdriver is not really an interesting tool, it will be your most-used tool. In addition to a standard screwdriver, you need a multi-bit screwdriver and a set of security bits, which have a small hole in their

center, allowing you to easily remove security screws that have small posts in their centers. These screws are another anti-tamper step that some manufacturers have introduced to help corporations reduce theft of expensive computer components, but they can cause a repair to come to a grinding halt as you try to figure out how to get into the case.

There are two common sizes of hex nuts on computers, so having both 3/16-inch and 1/4-inch nut drivers handy will make working with them much easier. You use 3/16-inch nuts for case standoffs (the small posts that support the motherboard and keep it off the surface of the case) and for the posts that mount most of the ports on the back of the computer, like monitors and printers. You use 1/4-inch nuts for expansion slots, power supplies, and case panels.

With all of the screws on computers, you would think that you would be able to find one when you need it, but that is not always the situation, so having a supply of standard screws (listed in Table 2-1) handy will be a great help to you. Since Compaq (and now HP) computers have standardized on the two hex screw sizes for optical drives and hard drives, and they use the hard drive screws for everything else. They have also been nice enough to usually provide a few spares for adding extra drives, which you will usually find screwed into the inside of the case somewhere.

Table 2-1	Stand Screws	
<i>Description</i>	<i>Thread</i>	<i>Length</i>
Hard Drive	6-32	5/32 in (4 mm)
Case	6-32	3/16 in (4.8 mm)
Optical and Floppy Drive	M3	1/4 in (6.35 mm)
Cable Plates and Ports	4-40	3/16 in (4.8 mm)
Case Fan	Self tapping	7/16 in (11 mm)

Miscellaneous tools

You need a variety of other hardware tools in your troubleshooting kit. This section presents a hodge-podge of miscellaneous tools of all sorts.

Digital infrared thermometers can find the temperature of CPUs, memory, and other internal devices without having to come into contact with them.

It is surprising how many times you need to read information on components that are already mounted inside the computer or need to see exactly where something has fallen. A small mirror can be a big help, and one that is mounted on an extension arm can be even better. This is useful when you want to see jumper settings or read hidden serial or model numbers. Again, for the cost, it is well worth the investment.

The inside of computer cases can be dark, and when you are working on a computer under a desk, the inside of the computer is even darker. Having a small flashlight in your set of tools is often more convenient than moving the computer to a location that has more light.

Some cases come apart easily after removing all of the screws, while others do not. The original Macintosh computer was one of those that did not come apart easily — by design. A *case cracker* lets you separate the pieces of a case that do not easily come apart, without causing the damage that a flat-head screwdriver would cause. This device comes in a range of sizes and shapes and is inserted in the gap of the case and then used to separate the sections. They were commonly used for Macintosh computers and monitors, which have a groove where the two halves of the case come apart.

Although vertical and horizontal hold isn't an issue with today's monitors, in the past, it was not uncommon to make adjustments to this element of the monitor. The adjustment knobs were usually located inside the case to keep users from playing with them and were reached with special tools. Monitor adjustment tools are shown amongst the tools in Figure 2-3 (from left to right: multi bit screwdriver with case crack tool, 3/16-inch hex nut bit, two monitor adjustment tools, wire cutter/stripper, diagonal cutter, extension pickup magnet, manual pickup tool, and extension mirror).

Figure 2-3: Several different tools that are good to have in your troubleshooting kit.



Loopback plugs serve two purposes, as a diagnostic tool and to eliminate errors messages. When you are testing parallel and serial ports, it is advantageous to have the system think that there is something connected to the port so that you can test sending and receiving data through the port. Rather than testing a serial port using a modem or null modem cable, and a remote computer to communicate with, you can use a loopback plug to test the port without any outside devices. Many multi-port gigabit network adapters ship with loopback plugs, which usually look like a short network cable, an inch or two long, with only one connector on it. This connector makes the network port think that it is plugged into a switch or hub and eliminates error messages from the operating system regarding errors with that port.

There are always cables and cable ties that need to be cut, stripped, or connected within the computer. Having a wire cutter, diagonal cutter, or wire stripper handy makes that job much easier. To go along with that, a knife or box cutter can sometimes be useful to make some detailed cuts. After your cutting is done, use electrical tape to wrap any exposed wires.

All slots should be covered, so you need to have a few spare slot covers with you for those cases in which you are removing a card from a system.

Similar to the slot covers, drive faceplates should be replaced when you remove a drive from a computer. These tend not to be standard across systems, so having spares that actually fit might be a little more challenging.

When mounting 3.5-inch drives in a system, you may choose to mount them in 5.25-inch bays, and to do so requires a mounting kit or adapter. Typically, this kit consists of a couple of side rails that make your drive wider.

Spare internal cables, especially the new 80-wire ATA-100/ATA-133 cables, are useful to have in your kit, as well as other standard internal and network cables. These help you diagnose cable or device issues. If you can replace a questionable cable with a known good cable, then you can rule out another element that might be causing the problem.

Network and other cable testers can be used to verify not only continuity, but also configuration of cables. In many cases, when people crimp their own network cables, errors can be made with pin configurations.

The “Hardware tools” section should be reviewed to a point where you can identify when you would want to use each type of tool.

Diagnostic software

In addition to various hardware tools that are available to help you troubleshoot a system, various software tools can help you out, too.



Boot/Rescue disk

To deal with many hard drive issues that might exist, you will want to have a boot disk at your disposal. Just booting your system with a Windows boot disk lets you determine whether you can access your hard drive. But to really be able to accomplish a troubleshooting task, or to recover a system, you need a little more power. More power can be found in many third-party boot disks and bootable CDs. Most of these solutions include a variety of testing and troubleshooting tools:

- ◆ Knoppix
- ◆ System Rescue
- ◆ WinInternals

Software diagnostic tools

There are diagnostic tools to test several major components in your computer, including, but not limited to, drives, processors, memory, serial and parallel ports, keyboards and mice, and network adapters. These testing tools typically verify integrity of components or stress components by performing multiple random or sequential reads and writes on the system.

BIOS and hard drive self-test

In addition to other software solutions, many BIOS routines include built-in testing software. These routines usually can test disks, RAM, processors, and other system components. Like the diagnostic software, these tests will usually perform random or sequential reads and writes to verify the integrity of the components that are being tested.



In addition to these tools, all of the tools listed in the Book VI, Chapter 3, which are designed for monitoring performance, can also be used for troubleshooting. Make sure you know the appropriate uses for each tool in both of these chapters.

The Art of Troubleshooting

In this section, you take a look at some general factors and then perform an overview of troubleshooting specific components. This chapter covers these components in the same manner as in the preventative maintenance chapter, Book IV, Chapter 1, by touching on each major component one at a time.

Troubleshooting basics

In the past, computer components were very expensive, but in today's market, most components have been turned into commodities and can be purchased very cheaply. Because so many elements are so cheap, replacing components is now more common than repairing them. With most components being inexpensive, it is easy to have a small supply of spare components and test components by swapping in new and reliable components.

Most failures in components occur either near installation or around the expected wear-out period; while very few components fail during the normal use period. With technology improvements proceeding at their current rates, it is very common for technology to become functionally obsolete prior to hitting the wear-out period for that device. This can be illustrated by CD and DVD drives over the last few years, where most CD drives that fail are going to be replaced with a CD-RW, CD-RW/DVD combo, or DVD-RW drive.

Physical environment

To troubleshoot and repair computer systems, you need a large, clean work surface and enough available power connections to power the equipment you are testing as well as your diagnostic tools. You should have anti-static straps and mats for working on equipment and a place to organize your tools nearby.

When doing remote repairs at a client site, you will want your travel tools, an anti-static wrist strap, and a clean space to work, which is sometimes difficult in a cluttered cube farm.

Audio and visual troubleshooting

After you have conducted your interview to get a list of symptoms, as suggested in Book I, Chapter 2, you will often start the first stage of active troubleshooting by looking for audio and visual cues that may be the cause of the described symptoms. This may come in the form of listening to *POST* (*Power-On Self-Test*) errors or beep codes or an examination of the physical environment.

POST errors

Each BIOS manufacturer has its own diagnostic codes that identify specific errors. You need to consult documentation for the specific beep codes for your BIOS. Many motherboard manufacturers will use codes similar to the original IBM POST codes, which are summarized in Table 2-2. If you get only one beep, all is good. In some cases, these beeps are also accompanied by a diagnostic code, which you also have to look up in the BIOS documentation.

<i>Beep Code</i>	<i>Description</i>
One short beep	Normal POST — system is okay
Two short beeps	POST error — error code shown on-screen
No beep	Power supply or system board problem
Continuous beep	Power supply, system board, or keyboard problem
Repeating short beeps	Power supply or system board problem
One long and one short beep	System board problem
One long and two short beeps	Display adapter problem (MDA, CGA)

Connectors and ventilation

As part of the visual troubleshooting, you will also likely check that intake and exhaust ports and vents are free of obstructions and that cables and connectors are securely attached. As part of the audio troubleshooting, you want to verify that you hear typical sounds coming from your computer, such as fans running properly and hard drives spinning. If the computer is too quiet, it may lead you to fan failures or other sources for your problems.

CMOS and BIOS

Modern CMOS settings are numerous and easy to access. If you have a computer that requires specific non-default settings to operate correctly, then you want to verify that default settings have not been loaded back onto the system. This can happen if certain jumpers have been moved on the motherboard, if the CMOS battery has failed, if firmware has been upgraded, or if there has been a strange failure within the system. Improper changes to the CMOS settings, such as boot order, video memory, power options, disk detection, swapping floppy letters, and CPU timing, can render the computer unbootable until the settings are corrected.

In some cases, you will have a problem with a specific BIOS or firmware version, and the motherboard or device vendor documents a fix through a newer version of the BIOS or firmware. If this is the case, refer to the update procedures and follow them. If they give you an option of backing up the current version of the BIOS, then you should perform that step as a recovery option.

Sometimes BIOS and firmware updates can cause more problems than they fix, and in some cases the update is one-way, and you won't be able to revert to older versions. After some upgrades, custom settings need to be reapplied.

Power-saving features in the BIOS may allow you to reduce the speed of the CPU to save power and reduce heat, and in some situations, this is used as thermal protection when the heat sink fan has failed. If you're looking at performance issues, this is one area to examine.

Motherboard

Motherboards are made up of many integrated components, and management of these components is done through CMOS settings or by changing jumper settings. You can find specific settings in your motherboard documentation. In addition, you can perform a visual inspection for damage, such as broken connections or capacitor damage, as shown in Figure 2-4. If you are looking for a failure in a component related to the motherboard, you should also verify that all relevant cables are connected to the motherboard correctly.

Processor/memory

Processor failures are rare, but not unheard of. In some cases, if it is immediately after installation, the issue may be related to seating of the processor. When the processor is added to the system, there will be a cooling mechanism as well, such as a heat sink and a fan. When you install the cooling mechanism, you should use either a thermal pad or thermal gel because it improves heat transfer between the processor and the heat sink. When there is fan or heat sink fan failure, many processors reduce the processor to a lower speed to allow for the generation of less heat.




Figure 2-4: This capacitor has sprung a leak. Note the residue on top.

Memory errors are often identified during the POST process, when memory is tested. It should be noted that soft reboots (using CTRL+ALT+DEL) usually skip POST tests, so full power cycles (using the power button) should be conducted from time to time. Fan and ventilation problems can also result in memory overheating and generating failures. Other errors occur when technicians unknowingly mix different types of memory, such as ECC and non-ECC memory (covered in Book II, Chapter 3), or memory that runs at different speeds.

Floppy drive

Many issues can affect floppy drives, but the most common is head alignment. When these drives cost hundreds of dollars, head misalignment was corrected by having a technician manually realign the heads, but because these drives are rarely used today — USB flash drives have become much more popular — it's usually cheaper to have faulty floppy drives replaced rather than repaired.

Other errors that can occur include damaged cables and connectors, improperly seated connectors for data or power, dirty heads, and failures in the media. Improper media selection can also be the culprit with some reported errors; there are still some single density and double density disks floating around with the high density disks. Some users with older computers may have used high density disks in their double density drives, making them unreadable in high density drives. Configuration settings in your CMOS (seen in Book II, Chapter 4), could also cause your floppy drive to not function.

Hard drives

When you're troubleshooting hard drives, don't overlook the physical problems. Many people dive straight to the software before checking all the hardware. If the drive is not detected or has data errors, you should check the connectors on the drive and motherboard and verify that the cable doesn't have any signs of damage. Sometimes just replacing the drive cable does the trick.

See if anything might have recently happened that could have caused a shock to the hard drive, such as dropping the computer. Other causes for errors and drive failure include heat problems, so verify drive placement, airflow, and heat dissipation for the drive. Other issues that can cause detection problems are the jumper settings for the drive. If the drive is set for Master, Slave, or Cable Select operation, then you could have problems if the setup is wrong for your situation. In some cases, you may have multiple jumper configurations for each setting, as shown in Figure 2-5, and all should be tried if there is a problem. For complete information about hard drive configuration, read through Book II, Chapter 5.

Figure 2-5: Some drives have multiple jumper configurations for the same setting.



If there are problems with disk size detection, then verify LBA (Large Block Access) mode if available or the use of drive overlay programs. Drive overlay programs are not used very much anymore because they add another layer of code to translate disk locations, which leads to slower access times, but they do let you use drives that are larger than what your system BIOS actually supports. Common disk size limits of some computer BIOS versions are summarized in Table 2-3. When possible, drive overlay programs should be avoided due to added complexity and reduced compatibility when working with that drive. Drive-mounted LEDs can be used to identify disk activity.

Table 2-3 Common Disk Size Limits

<i>Size</i>	<i>Limit Reason</i>
10MB	Early PC/XT limit.
16MB	Fat 12 limit.
32MB	DOS 3.x limit.
128MB	DOS 4.x limit.
528MB	Limit of CHS (cylinders, heads, sectors) mode, which has a limit of 1,024 cylinders. Fixed by LBA or INT13H translation.

(continued)

Table 2-3 (continued)

<i>Size</i>	<i>Limit Reason</i>
2.1GB	Limited by manufacturer-imposed LBA 4,095-cylinder limit or 22-bit LBA translation space.
4.2GB	CMOS extended CHS addressing limit of 8,191 cylinders, which was not widely used, or 23-bit LBA translation space.
8.4GB	BIOS-imposed INT13H 16,383 cylinder limit or 24-bit LBA translation space.
33.8GB	BIOS-imposed 26-bit LBA translation space.
136.9GB	Limit of ATA/66/100 BIOS 28-bit translation space.
150,994,944GB (144PB)	Limit of UDMA/133, ATA/133 “Big Drives 2001 specification” drives 48-bit translation space.



The ATA/ATAPI-6 or “Big Drives” specification was passed in 2001 and was incorporated into ATA/133 drives. This system was backward compatible with older drives and involved a change in the basic IDE circuitry. This also involved an introduction of the 80-wire, 40-pin ATA/133 cable. To ensure compatibility, the newer cables can be used for all drives.

On the software side, you can use `chkdsk.exe` or a similar tool on a rescue disk to test for file system and disk errors if the problem is with unreadable files or a non-accessible disk. If the problem is performance-based, then checking for disk errors is still a valid step, but you can also check for disk contention on the data bus and possibly move multiple drives to separate IDE/ATA channels or purchase an additional controller. Running Performance Monitor and monitoring disk counters will tell you if the disk is being over-used and is suffering from contention issues, which can be solved by moving some of the data files that are heavily accessed to an additional hard drive. To choose which data files to move to a new hard drive, you need to look at the actual applications that are running on the drive.

Other issues affecting drive performance include the type of bus that is being used — ATA 66 controllers and drives can be upgraded to ATA 100 or ATA 133 components; SATA 150 to SATA 300; and 5400 RPM drives can be replaced with 7200 RPM drives. When selecting a new drive, rotational speed, onboard cache, access time, and seek time should all be considered.

CD/DVD-ROM

As with hard drives, the connectors, cables, power, and jumper settings affect how the CD/DVD-ROM drives work. In addition, problems with optics and alignment can be a factor with optical drives, and a major shock can knock the drive mechanism out of alignment, requiring that you replace the drive. These drives also support analog audio, which makes an analog connection directly between your CD/DVD player and your sound card, allowing audio CDs to be played without requiring transfer of the data through your

ATA data cable. This analog audio cable needs to have its pin configuration verified at both ends of the cables.

Media problems can occur due to media storage or abuse. As I describe in the preceding chapter, some scratches may be repaired to a point that you can recover some or all data from damaged disks. Disk damage from glues, chemicals, or heat may be mostly or entirely unrecoverable.

Keyboard and mouse

Sometimes keyboard problems are easy to fix — the user enables scroll or number locks, which can be verified by the LEDs — but most of the time the problems are a result of damage from abuse or an accident, with the most common accident being spilled liquid. For the most part, these devices can be considered disposable because they cost so little to replace, but they still might require a few troubleshooting steps to identify where the problem lies. If the issue involves liquids or debris, such as food stuck under the keys, you might be able to recover this device with little effort. Food and other debris can be blown out with compressed air. If the liquid contains suspended particles (like sugar or salt) or is acidic, then the keyboard is likely a write-off, especially if the liquid has caused a short circuit or corrosion has set in. Cables may be damaged by keyboard trays, and both the cables and connectors should be checked for damage. For wireless devices, the batteries should be checked; failing batteries can cause a variety of usage problems.

If you have an urge to recover a keyboard that has liquid or debris damage, then you can soak it in deionized water, which is free of particles and is electrically and acidically neutral.




If you soak the keyboard in deionized water, the keyboard should be hung to dry for 24 to 48 hours. Don't apply heat to it to dry it faster. Adding heat can cause even further damage to the keyboard.

You can disassemble some keyboard models by removing the keycaps and the case screws (on the bottom of the keyboard). When disassembling a keyboard, remember that it contains a lot of little pieces, as shown in Figure 2-6, and they all have to go back in the correct places. Some keyboards also have springs under each keycap. After you have the keyboard apart, in addition to a bunch of small pieces for all of the keys, you'll also have a circuit board, which can be cleaned with an isopropyl alcohol-based cleaner and cotton swabs.

If the dirt is just on the keycaps, then keycaps can be removed with a special removal tool that has two wires that fit under the key and are then pulled directly up. After the keycaps are removed, they can be cleaned in a mild cleaning solution and replaced.

In general, keyboard repairs are not cost-effective due to the time it takes to make repairs.

Figure 2-6: Keyboards can be taken apart and cleaned, but they do contain a lot of little components.



Tracking problems are the most common ones that occur with mice; optical mice are fussy about the gloss level of the surfaces that they are used on, and internal rollers of mechanical mice can build up a layer of dirt. For an overview of maintenance of mice, refer to Book IV, Chapter 1.

Sound card/audio

Sound card and audio problems are often related to the connections to the external or internal speakers. Some sound cards and motherboards have additional connections for digital audio using an *S/PDIF* (*Sony/Philips Digital Interface Format*) connection, while others support 5.1- and 6.1-speaker surround sound. When using these types of connections, you should carefully verify the connections against the manufacturer's guide. When working with external speakers or surround sound systems, note that all of them require power to allow for amplification or separation of audio signals, so make sure that the power is plugged in.

Monitor/video

Some common and easy problems to fix with monitors involve user errors — adjusting the picture by using the knobs or buttons on the front of the monitor. There's nothing quite as annoying as going to a user's desk to look at a

“broken” monitor only to find that the user has reduced the contrast and brightness to zero, leaving the image faint or nonexistent. In addition to contrast and brightness, monitors usually offer horizontal and vertical size and position, curvature, and keystone adjustments.

When monitor problems are not related to these settings, they may be related to power input or video cable problems (which are easy to correct), or the internal power supply, electron gun, or picture tube. Because of the high-voltage capacitors and charges that are maintained by internal components like the CRT, repairs of the internal components should be left to a good repair shop.

Users have many common monitor symptoms. Here are some of the most common:

- ◆ **Incorrect colors** may be caused by front panel adjustments, bent or damaged pins on the video cable connectors (which can be easily fixed), or a damaged electron gun (which involves a trip to the repair shop or simply replacing the monitor).
- ◆ **Burnt or damaged pixels** are irreparable, so once the number is sufficient to impact the person using the monitor, it should be replaced or rotated to a user that is more tolerant of the missing pixels. I have had users extremely bothered by one burnt pixel in the top-left corner of their screen, while other users with a dozen random pixels that are burnt on their screen hardly even notice the dots are missing.
- ◆ **Blurry images** are usually an issue related to the monitor design and specification (like dot pitch) or a damaged or misaligned electron gun. In the instance of poor monitor design, you should not buy any more of that make or model of monitor, while, if you suspect a defect with the monitor, you should contact the manufacturer to see if the monitor can be replaced or repaired.
- ◆ **Screen flicker** may be related to the internal power supply but is more often related to trying to run the monitor above specifications for color depth, resolution, and refresh frequency; or the refresh frequency is just set too low. These are OS-level settings, so they can easily be resolved in the Display control panel.

Video problems can also be caused by your video card, which may not support settings to correctly communicate with your monitor and use appropriate resolution, color depth, and refresh settings for the monitor. If you have slow screen redraws, which can result in slow screen refresh or screen flicker, then the video card might not have enough processing power or video RAM. If your computer has an integrated video card, then you may be able to allocate more RAM to be used as video RAM. The final resolution for video card problems is getting your video card replaced.



All video problems that are reported by a user will have to do with the video card or the monitor or both, and will likely have to do with trying to get more out of the system than it is capable of.

Modem

With the proliferation of broadband Internet connections and companies being connected to the Internet, modem use has fallen off, but they are still utilized for special purposes. As with keyboards and other components, the cost of low-end modems has made the repair/replace evaluation lean toward replacement. Common modem problems include driver issues, line issues, and speed issues.

For several years, software modems, Win-modems, were popular because they had become so inexpensive compared to the old hardware modems. The migration from hardware modems to OS-level software modems meant that if the software drivers were not installed correctly, the modem would not work properly.

Phone line issues occur with your phone company's network and can often be identified by testing the phone line with another device or by using modem or OS diagnostic features (found in the modem's properties). A common line issue is call-waiting; when the tone interrupts the modem's signal. You can disable call waiting per call by adding a disable command to the beginning of the number you are dialing. Speed issues are often caused by improper configuration or by using both hardware and software compression (which can greatly reduce transfer speed).

Serial and parallel ports

Serial and parallel port problems can be caused by connectors, OS settings, and CMOS settings. Here are a few troubleshooting tips if you're having a problem with one of these ports:

- ◆ Check connectors for pin damage and check cables for damage at the ends or along the cables themselves. Often cables running along the floor get run over by users' office chairs.
- ◆ Both serial and parallel specifications have maximum limits to the overall cable length, and exceeding these limits can cause intermittent communication problems with devices. Parallel devices may support a standard Centronics interface, which has a cable length limit of 15 feet (5 meters), or they may support the new IEEE 1284, which has a cable length limit of 30 feet (10 meters). Serial (RS232) cables on the other hand have a length limit of 50 feet (15 meters) when communicating at 19,200 bps, but that limit increases to 3000 feet (914 meters) if you only communicate at 2400 bps.

- ◆ Serial and/or parallel ports might be disabled in the Device Manager. Right-click My Computer, choose Properties, click the Device Manager button on the Hardware tab of the System Properties dialog box, and make sure the ports are enabled. There's nothing worse than spending a few hours trying to install a printer only to find out that the parallel port was disabled.
- ◆ In addition to being disabled at the OS level, the port may be disabled in the CMOS, or in the case of parallel ports, you may be using a standard port when you need to support the bidirectional features of an ECP or EPP port. You can switch between SPP, ECP, and EPP standards for your printer port by changing your system CMOS settings.
- ◆ For both serial ports and modems, verify the connection speed and data characteristics, and verify that match with the device you are communicating with. These are specified by speed, data bits, parity, stop bits, and flow control.
 - The speed is measured in *bps* (*bits per second*) and ranges from 110 bps to 115,200 bps.
 - Speed used to be measured as *baud rate* (after its creator Jean Maurice Emile Baudot), which is a modulation rate or state change. Baud used to match up with bps (bits per second). With current modulation, encoding, and compression techniques, the bps rate is substantially higher than the phone line's baud rate (2,400 baud). Many people incorrectly use the term *baud* when referring to bps.
 - Data bits are either 7 bits or 8 bits.
 - Parity may be set to odd or even and is used to verify the contents of a byte (8 bits) of data. Parity uses the eighth bit of a byte and sets the total value of the byte to either odd or even, depending on the parity setting.
 - Stop bits are sent after the data to signify that the data byte has finished. Stop bits will be 1 or 2 bits long.
 - Flow control manages the rate of data transfer between the devices, and it allows for transfer at higher rates. Without flow control, the sender is able to send data only at a rate that guarantees that the receiver can deal with, or process, the data that it is receiving. When using flow control, if the receiver is about to be overwhelmed or needs a pause to clear a backlog of processing, it is able to use the agreed system to signal a pause or to request more data.



USB

Some common problems that you can encounter with USB include BIOS support, OS support, driver issues, version incompatibilities, and power requirements for devices. Here are some USB troubleshooting tips:

- ◆ As with all of the peripherals that I have covered, check all cables and connectors to ensure that they are properly connected and that the cables are not showing signs of damage or exceeding distance limitations for the technology.
- ◆ Not all system BIOS support or have adequate support for USB, so if you are encountering problems, check with the motherboard manufacturer for a firmware update that addresses your issue.
- ◆ Just like serial and parallel ports, USB ports can be disabled in the CMOS or through the OS, and this should be checked if the ports don't seem to be working.
- ◆ If the USB ports are enabled but are not working, then the issue could be hardware-related, and you can look at doing a repair or, in the case of a desktop computer, adding a separate USB controller via PCI or PCI Express.
- ◆ If the devices are detected but are not functioning, then you may want to ensure that correct drivers are available to the OS.
- ◆ Although most devices are backward-compatible, you may encounter performance or other errors when mixing USB 2.0 and 1.x devices and controllers.



USB 2.0 devices plugged into USB 1.x ports will perform slower due to the restrictions of the USB 1.x specification. If devices are not performing at the expected level, check for this situation.

- ◆ Finally, because most USB devices are powered from the bus, some devices may have power-related problems. While the USB specification provides up to 500mA (milliamps) of power to a single device, some USB controllers enforce the specification's low-power-mode startup at 100mA and then expect the device to request the necessary amount of power in 100mA increments.

If the device is dumb, such as a USB lamp, then it does not have circuitry to communicate with the USB controller. Since it cannot communicate with the controller, it will take whatever power it can get, so it may not get the power that it requires to function properly. Other devices, like some USB hard drives, require up to 1A (amp) of power, and this can be supplied by special double-connector USB cables or cables to draw power from other USB ports. Figure 2-7 shows the back of a USB hard drive, with a thick silver USB cable and a thin black power cable to be used for additional power. The connector on the end of the power cable is a pass-through connector, which would allow for another device to share the port that is being used to supply power to this drive. This sort of cable is required since some controllers will not allow devices to exceed the USB specified power per port.



Figure 2-7: Here is a device powered by two USB ports, using a pass-through connector for the second port.

All of the troubleshooting items for USB devices can also be applied to FireWire devices.

Power supplies and batteries

When troubleshooting power supplies, you first need to do the following two things:

- ◆ Make sure the power supply is configured for use in the country in which the computer lives. Different countries have different standards when it comes to supplying power, and the computer's power source needs to match that.
- ◆ Make sure the power supply is actually getting power. Checking this is easy enough — just see if the fan is running.

In addition, ensure that all power supply connectors are correctly attached to devices, including the motherboard. You can test the power coming from the wall receptacle with an inexpensive receptacle tester (recommended), with a non-contact voltage indicator (recommended), or with great care using a multimeter (not recommended due to the risk involved at the voltage

level being tested). Power cords can be tested for continuity by using a continuity tester or a multimeter.

Power supplies can be tested with inexpensive testers, as shown in Figure 2-8. The tester will ensure that the power supply is working and providing the proper voltages on each pin of the connectors. The tester in Figure 2-8 supports 20/24 pin power, 12V, peripheral, floppy drive, and SATA connectors. If you suspect that there is a problem with just one of the power connectors or one of the lines leading from the power supply, you can use a multimeter to test the connector, as described later in this section.

In some cases, your power supply may be functioning correctly, and your problem may be that you are attempting to power too many devices, and exceeding the total output of the power supply. Power supply testers will not verify that your total draw is not exceeding the amount supplied by the power supply; for that, you will need to resort to adding up the power usage by hand. For information regarding the draw by devices, see Book II, Chapter 6. If you suspect that you are exceeding the power output of the power supply, you can disconnect or unplug some of the devices and see if the system returns to normal operation (with the exception of the missing devices).

Figure 2-8: A power supply tester can save time when testing voltage supplied by power supplies.



If the power supply is not functioning or doesn't have enough juice to power all the components of your system, you should replace it. Other issues that are often a result of defective power supplies include:

- ◆ Power-up or system startup failures or lockups
- ◆ Intermittent rebooting or lockups during normal operation
- ◆ Intermittent memory errors
- ◆ Hard disk and fan simultaneously failing to run due to current shortage
- ◆ Thermal failures and overheating due to fan failure
- ◆ Electric shocks received when touching system case or connectors
- ◆ System operations halted or rebooting from static discharges

Since ATX power supplies use a soft power switch, they need to be connected to a motherboard in order to be turned on. So, in order to test the main power connector, the one that connects to a motherboard, you will either need to have the *power on* pin shorted (not recommended), use a power supply tester (recommended and preferred, refer to Figure 2-8), or back probing the connector when the computer is running (not preferred, but acceptable).




Power supply problems can result in problems with other systems. Check the question for keywords that suggest that there is an issue with the power supply, in addition to the other device(s) mentioned.

Most Molex-type power connectors can be *back probed*, which is done by having your multimeters black probe connected to ground and your multimeters red probe inserted into the back of the connector, which in most situations has sufficient space to allow the probe (see Figure 2-9). In this figure, the black probe is grounded, and the red probe is testing the 12V lead on the Molex peripheral connector. This result is valid as it is $\pm 5\%$, which is within the specification for power connectors. There is a minimal risk of damaging equipment if this procedure is done correctly, and it lets you see on the live system what the power issues might be. Refer to Book II, Chapter 6 to see what the appropriate voltages are for each pin.

Laptops

Many issues with laptops require sending them to an authorized repair depot, but there are many problems that you may be able to fix yourself. As a CompTIA A+ Certified Professional, you should also be able to verify which major component of the system is affected. But because laptops are highly integrated systems, many solutions to problems involve a main board or a motherboard replacement. Laptop manufacturers have also made it difficult to get into places that they do not want you to go.

Figure 2-9: Back probing power connectors is a standard method for testing power in systems.



Some standard issues that affect laptops include:

- ◆ Power-related issues with external power supplies, batteries, or capacitors
- ◆ External monitors or cutoff switches
- ◆ All of those Fn keys
- ◆ Pointer recalibration
- ◆ Tablet PC stylus and word recognition issues

When troubleshooting laptop power issues, you may want to start with the power supply or power adapter. Verify that it is the correct one for the laptop and that there is no damage to cables or connectors. You can check the power supply label for the positive lead and use your multimeter red probe on that lead with the black probe on the negative lead to verify the voltage output of the adapter. If the battery is not holding its charge, then it is likely just old and should be replaced — if it is new, see if it can be replaced under warranty. Less detectable is the fact that many laptops have power issues often dealing with overloaded capacitors or circuits. This can be solved by removing the battery, disconnecting the power, and pressing

the power button multiple times (as many as 20 times). Replacing the power cable should also fix this problem — if replacing the power cable doesn't work, it has at least given you a chance to test the system without the battery in it so that you can rule out a battery issue.

External monitors can also cause issues with laptops. On most laptops, when the external monitor is connected, it will be used instead of the laptop screen. When the monitor is turned off, the lack of a picture on the laptop can cause a call to you, so always check to see if an external monitor is plugged into the laptop. If this has happened, you can use your Fn function key to switch from the external monitor to the laptop screen. The specific key combination to switch between the external and laptop screens varies from one laptop model to another, but often it is the F4 key. You hold the Fn key down while pressing the F4 key at 2- to 3-second intervals to cycle through the display options.

Most laptops also have an LCD cutoff switch that turns off the display when the lid is closed, and it can usually be seen as either a small button at the top of the keyboard or a hole that has a mated post on the lid. Depending on the OS settings, triggering this switch may not only turn off the display but may also put the computer into standby mode or hibernation. For some models of laptops, this switch is susceptible to either being hit or stuck, thereby making the laptop think that the lid is closed.

Most laptops have three options built into a majority of their keys. Two of those functions are the same as on desktop keyboards and are activated with the Shift key, but laptops also have an Fn key and a series of extra functions identified by additional symbols in another color drawn on the keys. These functions include, but are not limited to, sleep, hibernation, brightness, contrast, battery meter, speaker volume, print screen, SysRq, scroll lock, num lock, pause, break, and all number pad functions.

To enable the number pad functions, which are shown in Figure 2-10, you do the same thing that you would do on a traditional keyboard: You enable the number lock. To do this on your laptop, you press Fn and NumLoc. After this is done, several keys on your keyboard will convert from their normal function to acting as a number keypad. These keys can be toggled for single key-strokes by using the Fn key as a Shift key. Often, when users accidentally enable the NumLoc, they call you because their keyboard is now “broken.” This at least should be an easy fix for you. Rather than dealing with this number pad, many users prefer to get a separate USB number pad, and there are some that do double duty as a standalone calculator.

Your laptop will have one or more pointing devices built into it. In most cases, it will be an Isopoint pointing device, which is a small post in the middle of the keyboard, or a glidepad or trackpad, which is a flat pad below the keyboard. Regardless of the type of pointing device that your computer has, it will sometimes need to self-calibrate. Self-calibration causes the

pointer to drift to one side of the screen and then to top or bottom. After it has finished, it will resume normal operation. Typically, this happens after a certain number of hours, and you may notice it once every few days. If it happens more often, then there may be a problem with the pointing device, and you may want to take the laptop to an authorized service depot. Many users will not know what is going on and will fight against the drifting and think that something is wrong with the system, but this is normal operation for these devices and thus a losing battle. As long as recalibration happens only periodically and takes only a few seconds to complete, the computer is working properly. Users should think of it as preventative maintenance.



Figure 2-10: Most laptops have a built-in number pad that you activate with the Fn key.

A tablet PC includes a stylus and word recognition software that converts handwriting to text. The stylus works with a special mesh layer on the laptop screen, and its positioning and tracking is controlled by an internal component called a *digitizer*. If there are problems with the digitizer, then it will be evident through stylus problems and will involve a trip to your authorized repair depot. Handwriting recognition is controlled through software, and on some systems may involve some training of the system to recognize your writing style. If there is a problem with recognition, then it may be necessary to re-train the system.

Getting an A+

In this chapter, you examine:

- ◆ Using a multimeter to test power supplies.
- ◆ Using antistatic mats and straps to reduce the chance of ESD damage.
- ◆ Third-party disks and tools are usually helpful when troubleshooting.
- ◆ Each major component inside the computer system has unique troubleshooting steps.

Prep Test

- 1** What is *not* typically measured by a multimeter?
 - A Current
 - B Capacity
 - C Voltage
 - D Continuity

- 2** What is the purpose of a network card loopback plug?
 - A To return signals from your network card back to your computer to verify accuracy and pin configuration of network cable.
 - B To allow use of 127.0.0.1 TCP/IP address range.
 - C To eliminate network error messages from the operating system.
 - D To capture network traffic for future analysis.

- 3** Which function is not usually associated with a laptop's Fn key?
 - A Number pad
 - B Immediate power off
 - C Monitor switching
 - D Sleep

- 4** Which issue doesn't indicate a possibly failing power supply?
 - A CPU thermal speed reduction
 - B Power-up failure
 - C System halt
 - D Internal devices receive 12V and 5V

- 5** Why would you use an extension magnet?
 - A To erase magnetic media in a drive.
 - B To suspend floppy drives inside metal cases.
 - C To pick up dropped metal objects.
 - D To remove Molex filings from inside of the connector.

- 6** What components can be tested by BIOS self tests? (Select all that apply.)
 - A Hard drives
 - B CPU temperature deviations
 - C Magnetic media
 - D Memory

7 How does the system BIOS report major startup errors or configuration issues?

- A Screen flashes
- B Beep errors
- C Error message dialog boxes
- D Audio interruptions

8 What is often the cause of reduced CPU performance?

- A Overuse of CPU cycles
- B Addition of RAM
- C Heat sink fan failure due to power supply failure
- D Removal of the J7 motherboard jumper as defined by the ATX 1.2 standard

9 What is the purpose of a security bit?

- A To reduce the chance of users opening system cases.
- B To secure TCP/IP data packets.
- C Eight of them create a security byte.
- D To allow for tracking of sensitive data.

Answers


- 1 B. Voltage, current, continuity, and resistance are commonly measured with a multimeter. *See "Multimeter."*
- 2 C. Network loopback plugs are used to reduce error messages related to unplugged network cables. *Review "Miscellaneous tools."*
- 3 B. The Fn key is not usually used to immediately power off the computer; that is usually done with the power button. *Check out "Laptops."*
- 4 D. Power levels for internal devices are supposed to be 12V and 5V. *Peruse "Power supplies and batteries."*
- 5 C. Extension magnets are used to pick up dropped metal items. *Take a look at "Miscellaneous tools."*
- 6 A, D. The system BIOS on some systems allow for testing of IRQs, memory, hard disks, CPU processing ability, and most internal components. *Peek at "CMOS and BIOS."*
- 7 B. Beep errors are the standard way that major system startup failures are reported to the user. *Look over "POST errors."*
- 8 C. Reduced CPU performance is often caused by the temperature of the CPU exceeding limits, and it reduces power to reduce the temperature. Excess temperatures usually occur when the heat sink fan fails. *Study "Processor/memory."*
- 9 A. Security bits are paired with security screws to prevent users from opening system cases. *Refer to "Miscellaneous tools."*

Book V

Operating System Basics

The 5th Wave

By Rich Tennant



Contents at a Glance

Chapter 1: Comparison of Major Operating Systems	499
Chapter 2: Operating System Functions.....	513
Chapter 3: Installing and Upgrading Windows	537
Chapter 4: Managing Files and Directories	579
Chapter 5: Command Prompt Procedures.....	613
Chapter 6: Working with System Files and the Boot Process	645

Chapter 1: Comparison of Major Operating Systems

Exam Objectives

- ✓ Identifying the features, functions, and hardware requirements of the Windows OS
- ✓ Identifying the features, functions, and hardware requirements of the Linux OS
- ✓ Identifying the features, functions, and hardware requirements of the Mac OS

Although it seems like we live in a Windows world, other operating systems live and thrive in the world. As a CompTIA A+ Certified Professional, you will encounter these operating systems from time to time. You may even find that some of these other operating systems are more capable than Windows for specific jobs that you need to do.

Of these other operating systems, the ones that you see the most often are Linux and Mac OS. In this chapter, I briefly introduce you to these operating systems, as well as to Windows, and discuss their capabilities.

What Is an Operating System?

All operating systems (OSes) are responsible for managing hardware, resources, and data. Each of the operating systems in this chapter is capable of performing these tasks with ease.

The responsibility of the OS regarding user interface consists of displaying an interface that always reacts the same way and makes it easy for users to perform the tasks that they want. This *GUI (Graphical User Interface)* attempts to make it easy and intuitive for the user to do what needs to be done. In an attempt to provide a space for applications to run and for the user to manage his or her environment, most operating systems with a GUI have employed the “desktop” metaphor. This metaphor declares the desktop as a work area on which you have a computer, a place to store items, and a waste bin to throw away stuff you don’t need or want anymore. All of these elements will also exist in the GUI.

The look and feel of the GUI, such as the look of the windows that are used, also applies to the applications that run on the OS. By having the applications use the same look and feel that the OS has and by having the OS provide common routines and dialog boxes for opening and saving files, there is a consistent appearance for the user, which usually reduces the amount of time required to learn new applications.



One thing that should probably be noted is that, although the hardware requirements vary and the installation procedures are very different, after you get into using the features of the OS, all three of these operating systems have similar usability features that allow a user with a short orientation period to become capable in any other OS.

Looking at the Windows OS

I start at the top of the OS hierarchy with the Windows OS. The look and feel of Windows today is very different from the look and feel of Windows 3.0 when I started using it in 1990. Windows XP and Windows Vista have streamlined the user interface and added features, such as Wireless Zero Configuration, to make it easier to use your computer.

Comparing versions of Windows

MS-DOS was awkward and intimidating for many people to use, as they sat there looking at the command prompt slowly blinking at them. Windows was first developed as a shell or user interface to go on top of MS-DOS, making MS-DOS easier to user. By the time Windows hit its third iteration, it was very usable as a tool. This product, Windows 3.0, went on to become Windows 95, Windows 95 OSR2, Windows 98, Windows 98 SE, and Windows ME. Windows ME represents the end of this product line, which has always been, due to architectural limitations, a shell on top of some level of MS-DOS.

The Windows that we now know and use owes its roots to a Microsoft/IBM venture that built upon IBM's OS/2 family to produce a new, secure server operating system. But in the manner of many partnerships, this one fell apart, and both parties took the fruits of their labors and went home. Microsoft took this work and created Windows NT 3.x. NT was short for New Technologies, and 3.1 was chosen to be compatible with the version number of the other Windows operating system. The Windows NT operating system was more stable than earlier versions of Windows because of its architecture (NT or *New Technology*) and more secure because of its file system (NTFS or *New Technology File System*). In addition to creating the server operating system, Microsoft created a desktop operating system in the form of Windows NT 3.x Workstation. Windows NT 3.x Workstation had the same core as the Windows NT 3.x Server, but had many of the server components removed.

The drawback to the operating system was that the architecture was totally different from that of the Windows 9x family. Because of this difference, many applications that were written for the Windows 9x family standards did not work with the Windows NT family. This limitation has slowly been overcome with each successive version of the operating system, sometimes at the cost of some of the stability features. This family went through Windows NT 4.0 to Windows 2000 and then to Windows XP. Windows XP was targeted not only to the existing users of the Windows NT family, but also to the users of the Windows 9x family, and as such, Windows XP took large steps at backward-compatibility with older applications. This focus on security, stability, compatibility, and usability of the Windows NT family continues with Windows Vista, which is the latest addition to the family.



The recent Windows operating systems of Windows XP and Windows Vista are based on Windows NT, and not Windows 3.0 or Windows 9x.

Typical system requirements


Although Windows started out with its roots in MS-DOS, the current versions of Windows that you will see are based on Windows NT Technology. For the most part, you will work with systems running on either Windows 2000 or Windows XP. Starting in January 2007, you will also see the next generation of Windows in the form of Windows Vista. Table 1-1 lists the system requirements for each of the versions of Windows.



Microsoft started using the term “Built on Windows NT Technology” with the release of Windows 2000 to tie the lineage of these two operating systems together. That would actually make the statement “Built on *New Technology* Technology.”

Table 1-1 System Requirements for Microsoft Windows

	<i>Windows 2000 Professional</i>	<i>Windows XP</i>	<i>Windows Vista</i>
Processor	133 MHz or higher Pentium-compatible processor	300 MHz or higher Intel- or AMD-compatible processor	800 MHz or higher modern processor
Memory	64MB or better	128MB or better	512MB or better
Hard Disk	2GB with 650MB of free space	1.5GB of available disk space	20GB with 15GB of free space
CPU Support	Single or Dual CPU	Single or Dual CPU	Single or Dual CPU
Display	VGA or better	Super VGA or better	DirectX 9.0 capable



Expect to get a question related to the hardware requirements for these Windows operating systems.

The biggest differences between the different operating systems are hardware requirements and some advancements in application compatibility. If you have a computer that is lower than or just meets the hardware requirements of an OS, you may want to consider using an older OS for that particular computer. For instance, although a computer with 512MB of RAM and an 800 MHz processor meets the requirements for Windows Vista, you might find that is like looking at a base model of a car, without all of its options. For any OS, if you choose hardware at the minimum level, you will get minimum performance. Using hardware above that level is like buying an options package for your car. At 512MB of RAM and an 800 MHz processor, you will likely find that Windows XP will offer somewhat better performance than Windows Vista, as it has less convenience and less overhead, and Windows 2000 would likely offer even more performance, but, in both cases, at the cost of losing the new features of Windows Vista. As you move backward in the Windows NT family, each version of the OS was designed to work with a lower level of resources. It then becomes a tradeoff of features versus speed. There is a minimum level of the OS that you will be able to work with, because of minimum requirements such as security or USB compatibility.

You are the only person who can determine which OS is right for you, and that decision depends on what features you actually feel are important and the hardware that you have available.

Identifying GUI elements

The Windows GUI follows the desktop metaphor, which places your files, active applications, and a virtual trash can all within easy reach on your computer screen. A taskbar at the bottom of screen allows you to quickly switch between open applications; it also holds a Start menu, which allows you to quickly launch applications that aren't currently running. All of this can be seen in Figure 1-1.

The Windows interface was introduced with Windows 95 and, ahem, *borrowed* the features that were working in the other major graphical OS of the time, the Mac OS. Although there have been some functional and cosmetic changes to the GUI, there have not been any major changes to how the GUI functions or operates.

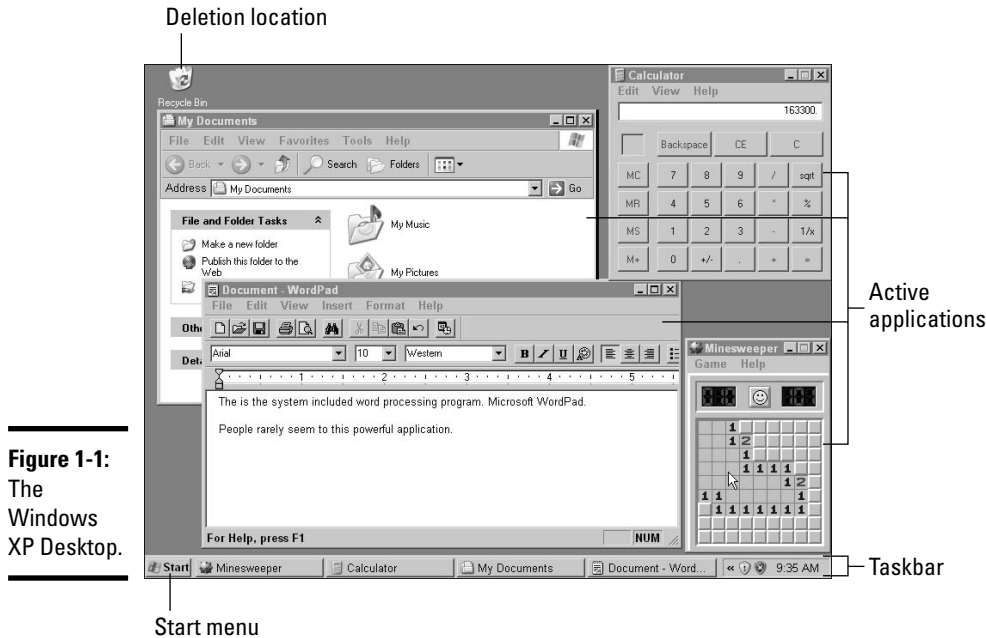


Figure 1-1:
The
Windows
XP Desktop.

Welcome to the Linux World

UNIX is an operating system that was originally designed to control telephone switches for Bell's telephone network. It was designed to be highly stable and configurable. As it grew to maturity, it became popular for its stability and processing power. Over the years, it went from solely a command-line environment to having a GUI interface. Its popularity led to a young university student's creation, during the early 1990s, of a UNIX-like shell that would work on Intel-based systems. This student was Linus Torvalds, and he created *Linus's UNIX*, or *Linux*.

Comparing versions of Linux

Since Linux's creation, it has been a worldwide group project, with Linus Torvalds controlling what components go into the Linux kernel, or the core components that make up the brains of the Linux OS. Many individuals and groups have contributed programming code to the Linux kernel, usually as completed drivers for hardware. The only thing that is common to all versions of Linux is the operating system kernel. Beyond the kernel, there are applications that are used with the kernel to give you a system that you can

actually use to accomplish tasks. If you wanted to, you could download the kernel source from <http://kernel.org>, compile the applications that you want to use, put them together, and you have a Linux system. Because it is that easy to make your own Linux system, there are currently just about as many Linux systems as there are tasks you want to accomplish with them.

When a group or company takes the Linux kernel, combines it with a user interface or shell, and then bundles in a variety of applications, they have created a distribution. All Linux distributions are cousins, and have been created using the same core components. Where distributions differ are in the mix of tools that are bundled in the distribution, and any customization the creator of the distribution decides to include with it.

There are a few major distributions of Linux, which have changed over the years. The current major distributions include:

- ◆ **Debian — GNU/Linux**
- ◆ **Fedora Core — Red Hat**
- ◆ **Gentoo**
- ◆ **Knoppix**
- ◆ **Slackware**
- ◆ **SUSE — OpenSUSE**
- ◆ **Ubuntu**

The main difference among these distributions is packaging. They all run the same kernel (or core), and most of the applications that they run are the same as well because the applications are written to the kernel specifications. So the real difference is the wrapping material that ties these elements together. In most cases, this wrapping includes an operating system installation tool and a packaging tool, which are used to prepare applications for installation and removal.

All distributions will have a command-line interface, and most also include a variety of GUIs. Common Linux-based GUIs include XWindows with several Window Managers, like AfterStep, Blackbox, Enlightenment, icewm, and Window Maker, to name a few. In addition to these standard Window Managers, Gnome and KDE are two complete desktop interfaces that are complete and popular.

Most of these distributions have packaging tools that are based on one of two standards — *Red Hat Package Manager (RPM)* or *Advanced Packaging Tool (APT)* — both of which offer the ability to install dependent packages

and resolve applications that conflict with each other. Both tools can also be configured to work with one or many networks or Internet package sources.

Applications that are written to Linux standards don't work on Windows-based systems but may work in the Macintosh OS X environment. Some developers write their Linux applications to common standards, such as Java or GTK, which both have components that can be installed on Windows systems.

Each Linux distribution has its own set of revision numbers to identify the exact version of the distribution that is being used. What matters more than the revision number is the version of the kernel that is used, as well as revision levels on the shared libraries, which are common routines that are used by many applications. Typically, when you install a new application, it has specific requirements for revisions of both the kernel and shared libraries but does not usually care about the distribution revision level.



In addition to the Linux distributions I mention earlier, BSD UNIX is making several new inroads into the market. BSD started in the 1970s and has spawned many child distributions, such as FreeBSD, OpenBSD, and PC-BSD.

Typical system requirements

Different Linux distributions have different system requirements. To take a distribution at random, SUSE Linux requires the following:

- ◆ 500 MHz modern processor or better
- ◆ 256MB of RAM or better
- ◆ 800MB of hard drive space
- ◆ 800 x 600 display

Most distributions have similar installation requirements and can install on substantially less hardware, depending on the purpose of the installation. For instance, some Linux distributions have been designed to run routers and firewalls, and these single-purpose distributions require substantially less hardware to function.

Linux runs on all of the same hardware on which Windows runs, as well as several hardware platforms on which Windows does not run. If you would like to test Linux, you can find several Live CDs that allow you to boot the distribution directly from the CD, without installing any files on your hard drive. These Live CD versions include Ubuntu, Knoppix, and others, which can be seen at www.linux.org/dist.



If you really want to get started with Linux, a great place to begin is *Knoppix For Dummies*, by Paul Sery (published by Wiley Publishing, Inc.). Not only do you get a *For Dummies* look at what Knoppix can do, but the book includes a copy of Knoppix 4.0 on an attached DVD-ROM!

A wide variety of file systems are supported by Linux; some are journal-based, like ext3 and reiser, and others are not, such as ext2. Journal-based file systems keep track of writes to the file system, and thereby provide better and faster recoverability in the event of a system crash or power loss. Most Linux distributions will install onto one of these file systems. As with Windows NT-based systems that support NTFS-based security, most Linux-based file systems support per user or per group security on files and folders.

Identifying GUI elements

Although many people use command-line systems exclusively when using Linux, there are many graphical user interfaces (GUIs) that can be used as well. The GUI loads as a separate element on top of the core command-line OS. Even when you are using a GUI, you will likely find that there are several occasions when you need to go back to the command line or a command-line window (called a terminal window) to accomplish your tasks. Figure 1-2 shows the Gnome desktop environment, which is one of the popular GUI environments for Linux. KDE is the other major desktop environment, and then there are several other Window Managers that may be used, such as Window Maker, IceWM, and FVWM95 (which looks like Windows 95), each giving you a different interface for window management on your Linux system.

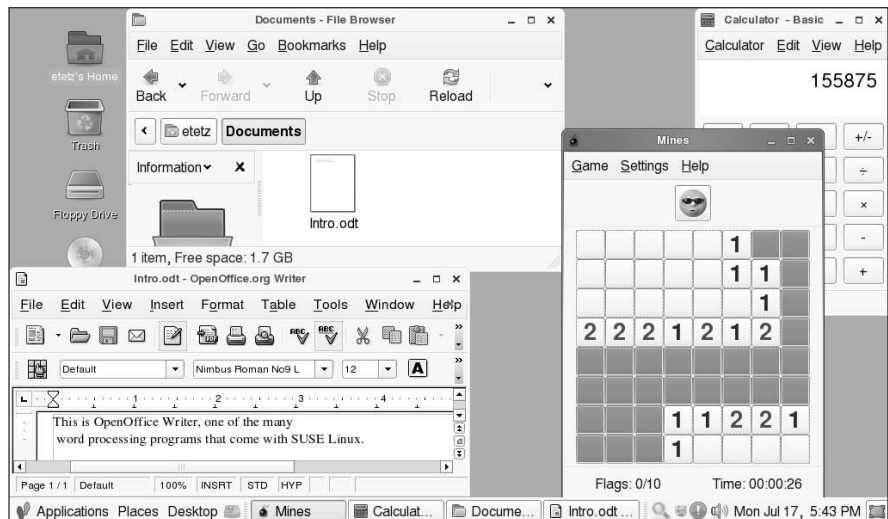


Figure 1-2:
The Gnome
desktop
environment
in action.

All of the GUI environments that are available for Linux have the same basic features for window management and work in conjunction with a file management application. The major differences between these systems are in the implementation and extra features, such as auto-mounting CDs as they are inserted into the drive. You should note that the Gnome desktop features and layout are similar to those used by Windows XP.

Understanding the Macintosh OS

The Macintosh OS has gone through some major changes over the last few years. The original version, Mac OS 1.0, remained essentially unchanged up through version 9.0. The last version of the Classic Mac OS was 9.x, and it was then replaced with a totally different OS, Mac OS X, which was based on the BSD OS.



The first nine revisions of the Mac OS were proprietary and built entirely by Apple, while Mac OS X and later are based on a BSD Unix core with Apple proprietary enhancements.

Comparing versions of Mac OS

The Mac OS originally hit the scene in 1984, running on a hardware platform that was very different from the hardware that was used by Windows PCs. It was standardized on the Motorola 68000 processor, while PCs were using Intel processors. The Macintosh also used small computer system interface (SCSI) hard drives, while Windows PCs were using integrated drive electronics (IDE) drives.

Through most of the 1990s, Mac OS 7.x was the main Macintosh OS. Shortly after the release of Mac OS 9.1 in 2001, Apple also released Mac OS X, with the goal of allowing capable computers to dual boot these two operating systems. Unlike Windows, Apple allowed you to dual boot the old OS to run any applications that did not work with Mac OS X, rather than having to support the old Mac OS 9 applications in Mac OS X. Mac OS X also offered to run old applications in an emulator called the Classic environment, which can be slow and requires that you have a Mac OS 9.1 system folder on your hard drive. Most Mac users were forced to quickly upgrade their applications to new versions or new applications that worked with the new operating system.

This change in operating systems and the break with the previous operating system allowed Apple to create a new, stable operating system platform for the Macintosh product line.

Typical system requirements

Unlike Windows, Linux, and UNIX — which can run on a wide variety of hardware, as long as it's IBM PC-compatible — Macintosh has always run on proprietary hardware with carefully managed BIOS and system ROMs. Until recently, these systems had always been powered by Motorola and IBM CPUs, starting with the 68000 family and more recently the Power PC G5 chip. Due to the processing requirements, Mac OS X requires at least a G3 processor to function.

Recently, Apple has switched over to using Intel-base CPUs in its product line — yet another component that is standard in PC hardware platform, blurring the line between the Macintosh hardware platform and the PC hardware platform.

Identifying GUI elements

Prior to Mac OS X, the Macintosh desktop environment was referred to as the Finder — the place where you would find all of the items that you needed to work with. There were no command-line interfaces for the OS because, unlike other operating systems, the Mac OS was developed from the beginning as a graphical OS.

With the adoption of the BSD-like kernel, the system has added a powerful new command-line interface based on UNIX to a powerful Apple developed GUI.

As shown in Figure 1-3, the same features that were present on the Windows and Linux operating systems' user interfaces are present on the Mac OS as well. One difference (or as a Mac user would think, one thing that is the same) is the Apple menu and the context-sensitive menu bar across the top of the screen. Unlike the other operating systems, which keep application menus with the application, Macintosh has always had the menu bar across the top of the screen. The Apple menu is always there (it's the equivalent of the Windows Start menu), and the remaining menu items change based on which application has the current focus.

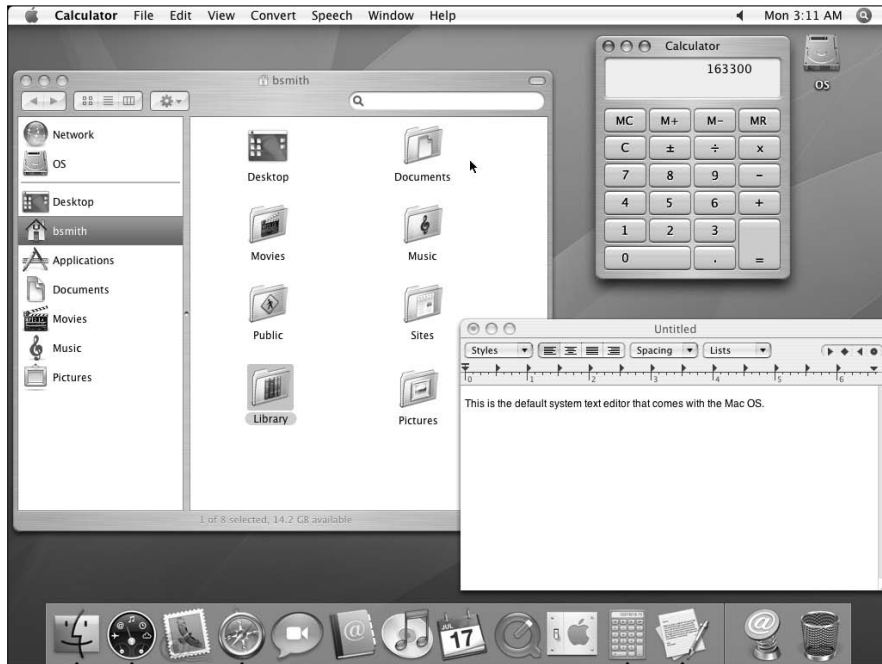


Figure 1-3:
The
Macintosh
GUI for
OS X.

Getting an A+

This chapter shows the differences between the major operating systems on the market — Windows, Linux, and Macintosh. Things to remember include:

- ◆ The major operating systems are Windows, Linux, and Mac OS.
- ◆ Windows and Linux hardware is compatible, but applications are not.
- ◆ Linux comes in sub-varieties called *distributions*.
- ◆ All three operating systems have similar GUI environments.

Prep Test

- 1** Mary works for a temporary placement agency and uses primarily Windows on her jobs. She is going to be sent to an office that uses Mac OS X. Should Mary expect to have large problems getting used to the new OS?
 - A Mary should expect to find the layout of the screen and the method of accessing applications to be completely different.
 - B Mary should refuse to visit the client site because the differences are so great.
 - C Mary should expect to find the layout of the screen and the method of accessing applications fairly similar to what she is familiar with.
 - D This is a moot point because the GUIs on both operating systems are identical.

- 2** You have been called by a client to give a recommendation on which OS should be installed on several new computers that are being added to the office. The other computers in the office currently use Windows-based clients (Windows 98) and Windows-based servers (Windows NT 4.0). What OS should you suggest?
 - A Red Hat Linux
 - B Free BSD UNIX
 - C Mac OS X
 - D Windows XP

- 3** Which of the following operating systems are direct ancestors of Windows XP? (Choose all that apply.)
 - A Windows 2000
 - B SUSE Linux
 - C Windows NT
 - D Windows ME

- 4** The variety of packaging of the Linux kernel and applications is called?
 - A Flavor
 - B RPM
 - C Revision
 - D Distribution

- 5** You just received a call from a user who is using a Linux workstation and has lost his graphical user interface. What is the likely cause?
- A** The user has exited the GUI and is now at a command shell.
 - B** The user has deleted the GUI, and it will need to be reinstalled.
 - C** The user has minimized the GUI shell window.
 - D** The user has a disk fragmentation issue that is preventing the GUI from functioning properly.

Answers

- 1 C. Mary should find that the interface is different, but not so different that she is unable to accomplish her tasks. If she wants to develop as high a level of competency as she possibly has with the Windows OS, then training would be beneficial, but not required. *See the "Identifying GUI elements" section of "Understanding the Macintosh OS."*
- 2 D. Since the office is already Windows-based and the users likely have a higher degree of proficiency with the new version of the Windows OS, a change to a more different OS should not be implemented unless the client has some specific goal that he wants to achieve by switching the OS. Even though, with the changes that have occurred with Windows XP, when compared to Windows 98 the user may be able to pick up any of the suggested operating systems with the same level of difficulty. *Review "What Is an Operating System?"*
- 3 A, C. Windows NT and Windows 2000 have a common heritage with Windows XP, in that they all use NT technology. Windows 9x operating systems, such as Windows ME, followed a different development path going back to Windows 3.0. *Check out "Comparing versions of Windows."*
- 4 D. A distribution includes packaging of the Linux kernel, libraries, and applications. It usually also includes a package of management tools and a system installation program. *Peruse "Comparing versions of Linux."*
- 5 A. Although you will need more information to properly diagnose the problem, when using Linux, it is possible to enter a system mode where the GUI is unloaded and you are left at a command line until you reload the GUI. *Take a look at the "Identifying GUI elements" section in "Welcome to the Linux World."*

Chapter 2: Operating System Functions

Exam Objectives

- ✓ Identifying the primary operating system components
- ✓ Describing features of operating system interfaces
- ✓ Using Remote Desktop Connection and Remote Assistance

Just as the engine is the driving force behind an automobile, the operating system (OS) is the driving force behind your computer. Your choice of engine in your automobile affects the performance of its system, and also affects what you can do with your vehicle. In a similar way, you will find that the performance and functionality of your computer is different if you use MS-DOS, Windows NT Workstation 4.0, Windows 2000, Windows XP, or some other operating system.

Because the A+ exam focuses on Windows 2000 and Windows XP, this chapter focuses on these two operating systems.

The operating system is responsible for several major functions on your computer, so you should take a look at some of these functions. In order to achieve its purpose of managing your computer, the operating system relies on its major components. This chapter examines these components and how they interrelate to effectively manage the computer. By the end of this chapter, you will be able to effectively choose between Windows 2000 and Windows XP for your system.

Identifying Major Operating System Functions

As you find out in the preceding chapter, the operating system is responsible for two major functions: managing hardware devices and providing an environment in which applications can execute. One of the major applications that needs to execute is the user interface, which lets the user control all other applications and their execution. In the case of Windows-based computers, this user interface, or *shell*, is named Windows Explorer. So together with hardware, this software makes up the operating system is responsible for getting things done.

The hardware and software get tied together in the location where everything happens: Memory.

Understanding memory management

Memory is the playing field where hardware and software mix. Memory management is the responsibility of the OS. The OS makes sure that the hardware and software components work within their own confines of memory. As memory resources are requested, the OS releases them either to the hardware device driver or to the application. The OS then takes steps to make sure that only the application accesses the memory areas that have been allocated to that particular application. If an application attempts to access memory that has not been allocated to it, the OS has to decide what to do with the application.

In all cases, attempts to access the memory space of other applications are denied. Usually the denied application doesn't know what to do when this happens, so you get an application that "hangs" or terminates through a *General Protection Fault (GPF)*. Most applications really never expect to be denied anything. Hangs and GPFs are annoying because the failed application is flushed from memory and data may be lost. On the positive side, the other application — the one whose memory was almost accessed — should still be working. The goal is to protect the other application by terminating the application that exceeded its boundaries. For more on GPFs, see Book VII, Chapter 2.

As part of memory management, the OS keeps track of physical memory on the computer as well as hard drive space that is used as extra memory available to the computer. Access time to physical memory is measured in nanoseconds, and access time to disk drives is measured in milliseconds. These scales are widely different, and so is the access time of information from these memory locations. Both types of memory are lumped together to make up virtual memory. Sometimes the term *virtual memory* is applied to the hard drive space that is used to simulate memory, but at the OS level, the term is used to refer to all of the memory available for storage on the system.

As applications launch and get loaded into memory, physical memory is used. When this space becomes limited, the OS moves some data out of RAM and onto the hard drive. The OS records the change in the actual location so that when the application requests that information again, the information can be moved back into physical memory and accessed.

Checking the OS version

No matter which OS you use, at some point, you will need to know its exact version number. This is often the case when you are considering installing

new software, which requires a specific version of the OS. There are many methods for retrieving this information in Windows 2000 and Windows XP.

The System Control Panel or System Properties lets you know which version of the OS you are actually running (see Figure 2-1). You can get to the System Properties dialog in the following ways:

- ◆ Choose My Computer→Control Panel→System.
- ◆ Choose Start→Control Panel→System.
- ◆ Right-click My Computer and select Properties.

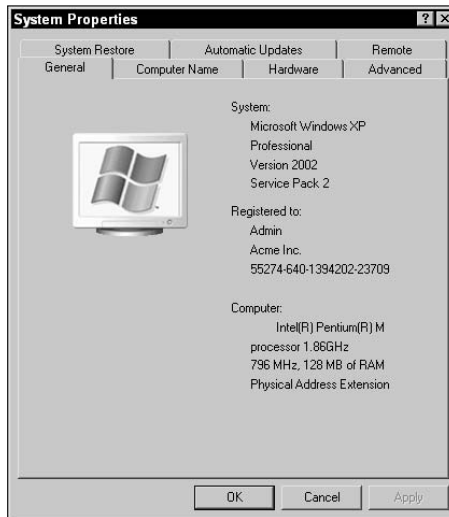


Figure 2-1:
The System Control Panel is available in all versions of Windows.

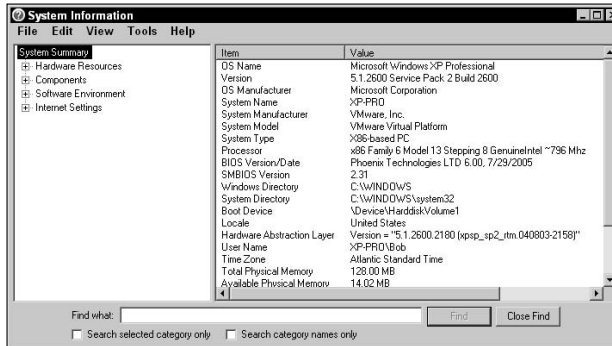
You see the version number of the OS listed in this System Properties dialog box, but Windows XP does not show the version information in the System Properties dialog box. For Windows XP and newer Windows OSes, you can use the System Information tool.

To launch the System Information tool:

- ◆ **In Windows XP**, choose Start→All Programs→Accessories→System Tools→System Information.
- ◆ **In Windows 2000**, choose Start→Programs→Accessories→System Tools→System Information.

The System Information tool includes the version number in the main window (see Figure 2-2). The System Information tool can either be launched through the Start menu, or you can use the Run command and type `winmsd`. `winmsd` (Windows NT Microsoft Diagnostics) is the name of the Windows NT 4.0 tool that would have provided the same information, but running it launches the System Information tool.

Figure 2-2:
The System Information tool for Windows provides a great deal of detailed information.



Some of the version numbers you might see are listed in Table 2-1. Notice that these version numbers are continuations of Windows NT version 4.0 in version and build numbers.

Table 2-1 OS Version Numbers for Modern Versions of Windows

<i>Operating System</i>	<i>Version Number</i>
Windows NT 4.0 Workstation	4.00.1381 Build 1381
Windows 2000 Professional	5.00.2195 Build 2195
Windows XP Professional	5.10.2600 Build 2600
Windows Server 2003 Enterprise	5.2.3790 Build 3790

Windows version numbers have three parts: the version number, the build number, and the Service Pack level, although not all tools will show you the Service Pack level. The version number (5.1 for Windows XP), and the build number (2600 for Windows XP) will not change for the lifetime of the product. As patches and upgrades become available for Windows XP, Microsoft releases hot fixes. *Hot fixes* are listed in the Add or Remove Programs Control Panel. When a sufficient number of fixes have been released, Microsoft releases a Service Pack, which incorporates most hot fixes as well as other unreleased fixes or features. The installation of a Service Pack adds a Service

Pack identifier to the end of a build number. The new version number looks like this:

Microsoft Windows XP Professional [5.10.2600 Service Pack 2 Build 2600]

After you know which version of the software you are running, you can take a look at what makes up your OS.

Understanding Major Operating System Components

In the following sections, you briefly examine how applications actually run on Windows computers. For a more complete description of how applications are supported, see Book VI, Chapter 2.

Getting into the architecture

Windows 2000 was a complete rewrite of the OS, and, as such, Microsoft was able to do things that it couldn't do with Windows 9x. When Windows boots, it immediately enters a 32-bit protected-mode state. The entire OS operates from this 32-bit state, and the kernel is loaded into *Ring 0*, which Microsoft refers to as *kernel mode* (see Figure 2-3). All processes running in kernel mode are protected from any processes running in *Ring 3*, or *user mode*.



The Intel processor design divides operation of the code into four separate execution levels, called rings. Ring 0 is at the middle of this arrangement, and Ring 3 is at the outside. These are the only two rings *that* Microsoft implemented during the development of 32-bit Windows operating systems.

The kernel is the core part of the OS that controls everything else that happens on the computer. The kernel is responsible for keeping user-mode processes separated from each other. Each application is started up in its own discrete area. One application is not directly allowed to interact with other applications and must pass such requests through *Executive Services*, which operate in kernel mode. This isolation of the applications from each other and from the rest of the OS is one of the keys to the stability of the current Windows system.

Kernel-mode processes are separated from user-mode processes but are vulnerable to corruption by other kernel-mode processes. The processor's architecture is all that protects kernel-mode processes from each other, and it often fails in its job. This means that Stop events (the ones that cause the infamous Blue Screen of Death), which reboot the system, are usually caused by a conflict of the processes that are running in kernel mode. For more on Stop events see Book VII, Chapter 3.

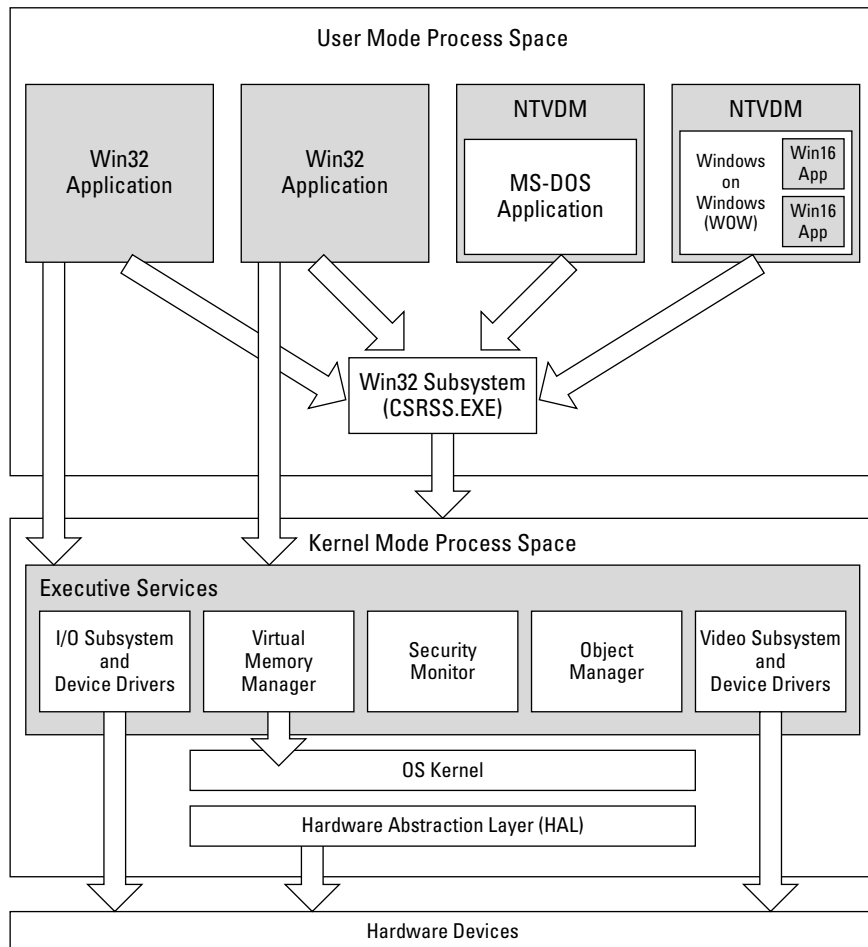


Figure 2-3:
The
Windows
XP archi-
tectural
overview.

Paging your memory

Windows uses hard drive space to extend the amount of memory that is available to applications. The total pool of memory that is available is referred to as the *Virtual Memory Page Pool*, or just *virtual memory*. Virtual memory is composed of both physical RAM and hard drive space and is maintained by the *Virtual Memory Manager (VMM)*. Figure 2-4 shows how these systems relate to each other.

The Virtual Memory Manager is responsible for keeping track of where different kinds of information are located in memory. As applications request to have information placed in memory, the VMM places the information in an

area of RAM. If RAM is getting full, then some information is moved from RAM to an area on the hard drive. The location that is used on the hard drive is either called a *swap file* (in older version of Windows) or a *paging* or *page file* (in current versions of Windows). As applications request stored data, then the VMM moves information around so that the data is available in RAM.

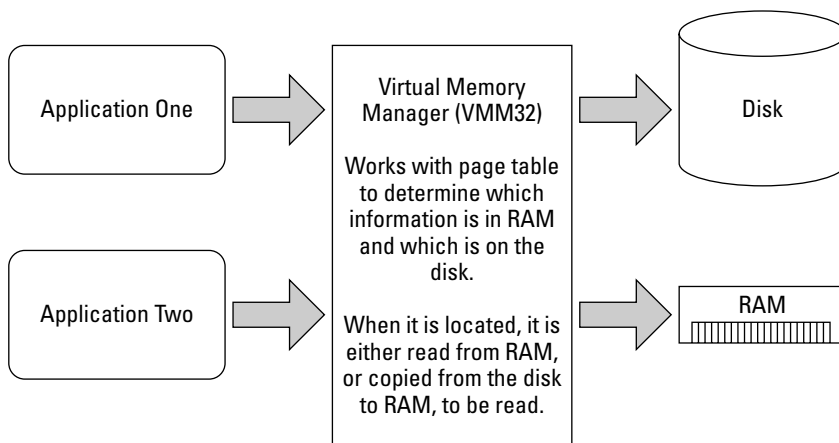


Figure 2-4: A sample of virtual memory within the Windows XP OS.

You configure the paging file in the System Properties. Follow these steps to adjust how much virtual memory your computer uses:

1. Open the System Properties dialog box.

Windows offers a few routes to get there:

- Right-click the My Computer icon on the desktop and choose Properties.
- Right-click the My Computer icon in the Start menu and choose Properties.
- Choose Start → Control Panel → System.

2. In the System Properties dialog box, click the Advanced tab.

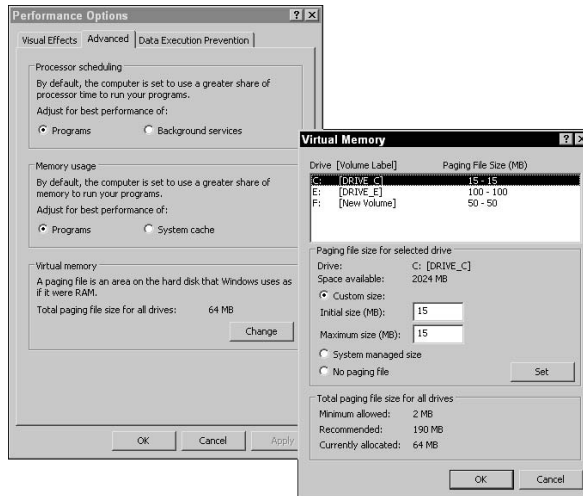
3. Click the Performance Settings button.

The Performance Options dialog box appears.

4. In the Advanced tab, click the Virtual Memory Change button (see Figure 2-5).

The Virtual Memory dialog box appears. In this dialog box, you can set an initial and a maximum size for the paging file on each logical drive.

Figure 2-5:
Page file settings for Windows XP can be used to improve performance.



5. Adjust the settings as needed.

You might want to adjust the virtual memory settings for the following reasons:

- **You are running out of space on your boot partition (the one with the Windows directory).**
- **You want to improve paging performance by reducing disk contention with the OS.** In this case, you can move the paging file to a different physical disk if you have one.
- **You want to improve paging performance by load balancing the page file between different physical drives.** Unlike Windows 9x, current Windows systems allow you to have multiple paging files, each on a different disk.



If you reduce the size of the paging file on your boot partition below the size of the physical RAM on your system, you are warned that some of the recovery options will be disabled. The paging file is used as a storage space for the `memory.dmp` file that is generated during a Stop error.

6. Click OK on three consecutive dialog boxes to close them all out.

Choosing your file systems

Just like many companies have different systems that for filing data, with Windows, you have several options when choosing a file system. MS-DOS gave us the File Allocation Table (FAT), which started out as FAT12 and then became FAT16 in later version of MS-DOS. Windows 95 then added to that

base file system and gave us FAT32. In between this time, the different architecture of Windows NT gave us the New Technology File System (NTFS). The latest versions of Windows, like Windows 2000 and Windows XP, support all three file systems, and you get some control over how to use them.

When using Windows XP, your choices include FAT, FAT32, and NTFS. Table 2-2 summarizes some of the differences between the file systems.

Table 2-2 Differences between the Major Windows File Systems

<i>File System</i>	<i>OS Support</i>	<i>Partition Size</i>	<i>Max. File Size</i>
NTFS	All versions of Windows NT (including Windows 2000 and Windows XP) support NTFS, but there were major changes to the file system with the release of Windows 2000. Computers running older versions of Windows NT may not be able to access the files.	10MB up to a theoretical limit of 16EB.*	Limited only by the size of the volume.
FAT32	FAT32 was released with Windows 95 OSR2, and all version of Windows 9x after that support FAT32. Windows 2000 was the first OS in the Windows NT family to support FAT32.	Up to 2TB,* usually used for FAT volumes larger than 512MB. Windows 2000 and XP will not format partitions larger than 32GB with FAT32.	4GB
FAT	Accessible by all Microsoft operating systems since MS-DOS, including Windows, Windows 9x, and Windows NT-based operating systems. It is also widely supported by other operating systems.	Up to 4GB. MS-DOS and Windows 9x allow you to create FAT partitions only up to 2GB, while Windows NT-based systems can create partitions up to 4GB.	2GB

* One terabyte (TB) equals approximately 1,000 gigabytes. One exabyte (EB) equals approximately 1 billion gigabytes, or 1 million terabytes. That's a lot of memory!

In most cases, you want to use NTFS because:

- ◆ **NTFS supports for security on folders and files.**
- ◆ **NTFS is a journaling file system, so it offers better recovery in case of power interruptions.**
- ◆ **NTFS offers a more stable platform for data storage.**



Although some older applications may have compatibility issues with NTFS, these issues are rare. For more on file systems, see Book II, Chapter 5.

NTFS should always be chosen as the file system due to stability and security. You will only choose FAT or FAT32 if you need to dual boot a computer with another operating system, such as Windows 9x or Linux.

Registry

Windows uses its Registry to record settings for applications and for the operating system.

When people first started writing applications for Microsoft-driven computers, they usually needed to record a series of per-use or per-computer settings. In most cases, they wrote these settings to a file stored on that computer. This system worked fine most of the time, but the people who needed to manage those computers and applications found that some of these configuration files were binary files and some were text files, and occasionally they would get corrupted and be unusable. To solve some of this confusion, Microsoft came up with a file standard that appeared when Windows was introduced to the world.

The new solution was the `.ini` (“innie”) file, a text file that uses sections, settings, and values stored in a standardized format. Because the data is written and can be retrieved using a standard format, the operating system can use standard procedures to store and retrieve the data rather than relying on each application developer to write their own procedures. Listing 2-1 shows a portion of a sample `.ini` file.

Listing 2-1: A Sample `.ini` File

```
[Options]
WordSel=0
Units=0
Maximized=0
FrameRect=AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
PageMargin=AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
DefaultFormat=5
[Text]
Layout2=CAAAAAAAAAEAAAAAAAAADAAAAAAAA
LayoutAux2=CAAAAAAAAAEAAAAAAAAADAAAAAAAA
Wrap=0
BarState0=0
BarState1=0
[RTF]
Layout2=CAAAAAAAAAEAAAAAAAAADAAAAAAAA
LayoutAux2=CAAAAAAAAAEAAAAAAAAADAAAAAAAA
```

```
Wrap=0
BarState0=0
BarState1=0
```

You should note that the section names are enclosed in square brackets.

Originally, these `.ini` files could be stored in the application directory, in the Windows directory, or anywhere on the file system path, and this location variability often led to problems, such as multiple, conflicting `.ini` files and uncertainty as to which file was in use.

With the release of Windows 95 and Windows NT, Microsoft programmers developed a new way to deal with this issue — they created a single location in which to store all configuration information for a computer. This was called the *Registry*. Like `.ini` files, the Registry uses a standard format. The Registry breaks settings down into two basic categories: settings for computers and settings for users. Computer settings are stored in a section named `HKEY_LOCAL_MACHINE`, while user settings are stored in `HKEY_CURRENT_USER`. Application developers are encouraged to use this new location to store settings for their applications and have even been given a special software key to store their settings. It is up to the software developer to decide whether settings are user- or computer-based and to manage the settings appropriately.

Even though Microsoft provided the Registry as a location to store program settings, not all programmers use it for their programs, and some still use `.ini` files. You should expect to encounter some programs that are still using `.ini` files.



Remember that Registry stores all settings in either `HKEY_LOCAL_MACHINE` or `HKEY_CURRENT_USER`, depending on whether the setting is related to the computer or the user.

For detailed information about how to view and manage settings stored in the Registry, refer to Book VI, Chapter 4.

Navigating Your Computer

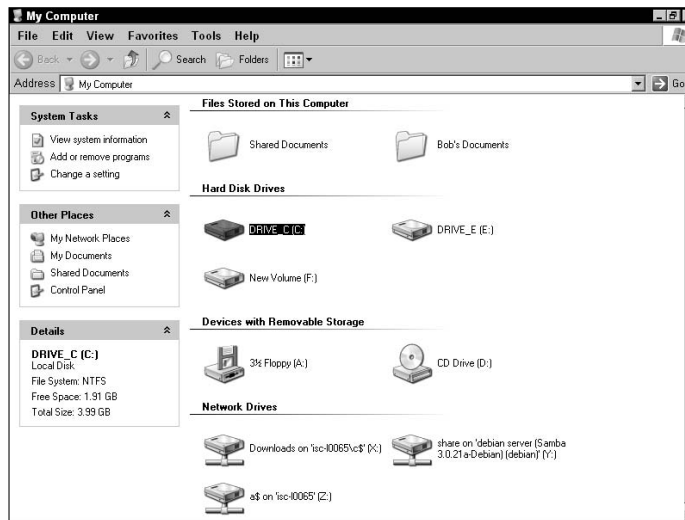
Being able to get to the files you need when you need them is crucial. From the command prompt, you can work with `cd` (change directory) and `dir` (directory), which are covered at length in Book V, Chapter 5. When you are using Windows, you can choose between Windows Explorer (not Internet Explorer) and My Computer. Either tool can be used for navigation, and the reason for choosing one or the other will be a matter of personal preference. This section will look at these tools in a fresh or default Windows XP installation using the

default settings, if you have changed your settings, then some windows may not be the same as I am describing. When using these tools, Windows XP has a few more options than Windows 2000.

My Computer

My Computer is the most common way people start accessing files on their computers because it's easy to find in the top-left corner of the desktop. When you double-click My Computer, a large window is displayed (as shown in Figure 2-6), listing the drives on your computer. You can then double-click your C: drive to see what files are located on that drive.

Figure 2-6:
My
Computer
starts
navigating
through the
file systems
at the drive.



The toolbar at the top the My Computer window features navigation buttons that allow you to go back to the previous folder or up to the parent folder. There are also a set of Cut, Copy, and Paste buttons that allow you to move files between folders.



Several tasks are more easily accomplished with an additional open window, such as copying files from one folder to another. To open an additional window, you can either double-click My Computer on the desktop and browse to the correct folder or hold down your Ctrl key as you double-click a folder, which leaves the original window in place and opens a new window for the folder that you have double-clicked on.

If you don't like how the files in the My Computer window are displayed, you can use the Views button on the toolbar to change the view from Icons to

Thumbnails, Tiles, List, or Details. Thumbnails will let you see the contents of many files, such as graphics, while Details view offers the advantage of listing file sizes and modification dates, which is often useful when trying to locate specific files. The choice of viewing style is really a personal preference.

The address line near the top of the My Computer window allows you to type the name of the folder that you want to go to and directly switch to it. For example, you could type `C:\Windows` to switch to the Windows folder.

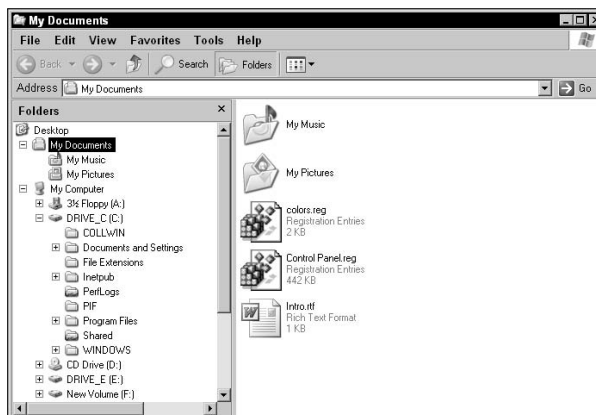
When you select an item in the My Computer window, you see properties for that item in a panel on the left side of the window, as well as common tasks that can be performed on the selected item. This is dependent on the window being large enough to display the information for you.

If you want to explore many of the options for displaying information in the My Computer window, then open the Folder Options dialog box from the Tools menu.

Windows Explorer

By the end of this section, you will see that My Computer as an application does not actually exist. To open this application, choose `Start→All Programs→Accessories→Windows Explorer` for Windows XP computers, or `Start→Programs→Accessories→Windows Explorer` for Windows 2000 computers. This opens a window like the one shown in Figure 2-7. This window has a right-hand pane that resembles My Computer and a left-hand pane that hierarchically displays all folders or containers on your computer, starting with Desktop and defaulting to My Documents. The left-hand pane is called the *Navigation pane* or *Folders Explorer Bar*. The left-hand pane is what makes Windows Explorer different from My Computer.

Figure 2-7:
The hierarchical view of the folder structure makes Windows Explorer a very useful application.



Now, you will see why a My Computer application does not really exist. In the top-right corner of the Navigation pane is a small x, which is the close button for the navigation pane. If you select My Computer in the Navigation pane, close the navigation pane, and change the view to Icons for both Windows, you will see that this window is the same as the My Computer window. To show the Navigation pane, select View⇨Explorer Bar⇨Folders. The Navigation pane allows you to quickly move from one folder to another, or to copy files from one folder to another without opening multiple windows. Many Windows users prefer the Explorer view to the My Computer view of their files.

This section showed that both views of your files use the same program, Windows Explorer or `explorer.exe`.



`explorer.exe` is not only used for My Computer and Windows Explorer, but the first time it is run on your computer, it is launched as the desktop. If you close all open windows and check your Processes tab in Task Manager, you will see that one copy `explorer.exe` still running. If you right-click⇨End Process on `explorer.exe` in the processes list, your desktop will disappear, and you will only see your wallpaper. If `explorer.exe` does not automatically reload, use Task Manager, File⇨New Task(Run) to open the Create New Task dialog box, type `explorer.exe`, and click OK. This will reload the desktop copy of `explorer.exe`.



If you want to quickly open Windows Explorer, right-click My Computer and choose Explore from the shortcut menu. This Explore option shows up in the right-click menu for any folder.

My Network Places

My Computer is great when you're working with data that is stored on your computer, but, in the connected world that we live in today, more and more of the data that you need to work with is stored elsewhere, across a computer network. My Network Places (or Network Neighborhood in earlier versions of Windows) is the most convenient way to find those network resources that you need to work with. You will find My Network Places on your desktop and in your Start menu.

Regularly visited network locations can be stored in My Network Places, making them easy to locate when returning to those locations. To add items to this list, click on the Add a network place, which is in the Network Tasks section of the left-hand pane, to open the Add a Network Place Wizard and then step through the wizard's instructions.

Command line

The last way to get around on your computer is the oldest of the methods: the command line. Some people refer to this as MS-DOS, but when you're dealing with Windows, you're normally running the 32-bit command-line application. This application is capable of executing any command found on the system, but is most often used for commands that do not have a Windows GUI component. Microsoft has command-line tools available that allow you to perform most system management tasks from the command line, which is useful when troubleshooting or when creating scripts or batch files to perform tasks.

You access the command line in Windows XP by choosing Start→All Programs→Accessories→Command Prompt. In Windows 2000, choose Start→Programs→Accessories→Command Prompt to open the Command Prompt window.

Using Tools and Configuration Utilities

A computer's configuration is very important for both cosmetic and functional reasons. In the following sections, you get a good look at the tools that you use to manage your system configuration.

Control panel applets

Most of the settings for your computer are stored in the Registry, but the Registry is a place where, if you make an error while making changes, it can require that you reinstall your operating system. To make your actions less prone to major errors, Microsoft created the Control Panel. The Control Panel contains applets that let you change many system settings without requiring you to make direct changes to the Registry. Control panel applets are the recommended method of changing most of your system settings. Table 2-3 provides a summary of each Control Panel applet and identifies on which operating system(s) you'll find them.

Control Panel Applet	Description
Accessibility Options	Changes the settings for accessibility features for users who have physical limitations.
Add Hardware (XP)	Runs a Plug and Play hardware detection of your computer to detect new hardware devices or allows you to manually select devices to install.
Add/Remove Hardware (2000)	Performs Plug and Play detection, full hardware scan, device configuration, and device removal.

(continued)

Table 2-3 (continued)

Control Panel Applet	Description
Add or Remove Programs	Allows you to start the installation of programs, displays a list of installed programs that can be removed, and allows you to add or remove optional operating system components (such as games).
Administrative Tools	Opens a folder that contains the Administrative Tools for Windows. These tools allow you to configure, manage, and control your Windows computer.
Automatic Updates	Configures the Automatic Update service on Windows and allows you to choose how and when your updates are installed.
Date and Time	Changes the system date, time, and time zone.
Display	Allows you to configure your display options. This includes video card settings, screen resolution, color depth (number of colors), and color scheme.
Folder Options	Allows you to change the default folder options for your computer. These are the same options that are available through Windows Explorer's Tools⇨Folder Options command.
Fonts	Displays a list of fonts that are installed on your computer, allows you to see what the fonts look like, and allows you to delete fonts from or copy fonts into the folder.
Game Controllers	Allows you to add or remove game controllers and joysticks.
Internet Options	Allows you to configure Internet options and connection settings. It also allows you to clear your Temporary Internet Files and History folders.
Keyboard	Allows you to change the sensitivity settings for your keyboard and the layout for your keyboard.
Licensing (Server)	This is available only on Windows servers. It allows you to set or change the number of client access licenses that you have purchased.
Mail	Configures the mail system for the computer.
Mouse	Configures the orientation, motion, and double-click settings for your mouse.
Network Connections (XP)	
Network and Dial-up Connections (2000)	Configures connections to your network or other networks.
Network Setup Wizard	Sets up your computer to run on a small network. Often used for home environments.
Phone and Modem Options	Configures dialing rules and modem settings.
Power Options	Configures power management settings for your computer.

<i>Control Panel Applet</i>	<i>Description</i>
Printers and Faxes (XP)	
Printers (2000)	Adds, removes, or configures printer settings and queues.
Regional and Language Options (XP)	
Regional Options (2000)	Configures the operating system to support different currency, numeric, date, and time settings for specific countries, and inputs and displays languages.
Scanners and Cameras	Configures installed scanners and cameras.
Scheduled Tasks	Lets you add, modify, or delete scheduled tasks.
Security Center	Configures and reviews security settings for the computer, including anti-virus, Windows firewall, and automatic updates.
Sounds and Audio Devices (XP)	
Sounds and Multimedia (2000)	Lets you choose sounds for system events, like window open or new mail, and allows you to configure audio or multimedia hardware.
Speech	Lets you change settings for text-to-speech systems or for voice recognition.
System	Lets you configure advanced system settings for hardware and performance.
Taskbar and Start Menu (XP)	Configures settings for items displayed in the Taskbar and in the Start menu.
User Accounts (XP)	
Users and Passwords (2000)	Configures user profiles and user settings.
Windows Firewall	Contains configuration settings for the Windows Firewall that comes with Windows XP Service Pack 2.
Windows Media Connect	Configures digital media devices.
Wireless Link	Configures infrared ports on system to send and receive files.
Wireless Network Setup Wizard	Wizard for setting up or adding to a wireless network at a small office or home.

Microsoft Management Console

One common complaint about earlier version of Windows was that you had to use a different management tool for every task. For example, Windows NT 4.0 has tools for managing users or computers in a network domain *and* locally on the computer, tools for managing performance, and tools for managing disks and disk partitions. With Windows 2000, Microsoft took steps to consolidate these disparate tools; their solution was the *Microsoft Management Console (MMC)*.

The MMC is a framework into which other tools, or *snap-ins*, can be loaded. All of the items that you see in your Administrative Tools folder represent MMC sessions configured with preset snap-ins. To see the raw interface, run `mmc .exe` from the Run command or from the command line.

Taskbar and Notification Area

To let you keep an eye on what programs are running on your computer and to allow you to switch quickly between running programs, Microsoft created the Taskbar. The Taskbar is movable, but by default it is at the bottom of your desktop. One end of it has the Start menu, the other end has the System Tray, or Systray, and the space in between is made up of tiles representing all of the current running applications. Starting with Windows XP, Microsoft calls the System Tray the *Notification Area*. If you find an empty area of the Taskbar, you can right-click and choose Properties. You will be able to modify settings for both the Taskbar and the Notification Area, such as auto-hiding the Taskbar, displaying the Quick Launch toolbar, grouping buttons, showing the clock, and hiding inactive icons.

Start menu

If all of the items on your desktop were missing, then you would still be able to accomplish all that you need to do by using the Start menu. The Windows 2000 Start menu was always very functional, but Windows XP introduced enhancements that put more resources, like My Computer, My Documents, and My Network Places, right at your fingertips.

The Start menu, like its name suggests, can be used to start most of the applications that are installed on your computer. Contrary to its name, you can also use it to shut down your computer and stop everything.

Remote Desktop Connection

Even though you can accomplish a wide variety of tasks at your computer, sometimes you'll want to do something on other computers on your network. Sometimes, you'll just want to modify files that are found on the other computer via shared folders, while other times you will need to actually launch and run applications on that computer. Remote Desktop Connection allows you to do this.

Microsoft has invested time creating *Remote Desktop Protocol (RDP)*, which defines how to send key strokes and mouse movements across a network from one computer to another computer, and how to have what is displayed on the screen of a computer sent to another computer. Using these mechanisms, you are able to remotely use another computer on your network, as if you were sitting in front of that computer.

Remote Desktop Connection allows you to connect across a network to another computer and remotely take over the desktop or console by communicating with it using the *Remote Desktop Protocol (RDP)*, as shown in Figure 2-8. From a computer running Remote Desktop Connection, you can connect to a *Virtual Desktop* on a server running Windows 2000 Server or Windows Server 2003, both of which can support hundreds of remote users simultaneously, or you can connect to the desktop of a Windows XP Professional computer. In either case, this allows you to run applications on the remote computer as if you were sitting at the computer. Your computer passes mouse movements and keystrokes up to the remote computer, and you see virtual screen shots of what is happening on the remote computer screen. These are virtual screen shots because what you are seeing does not appear on the remote computer. You will find Remote Desktop Connection in the Start menu under All Programs→Accessories→Communications, or you can launch it from the Run command by typing `mstsc.exe`.

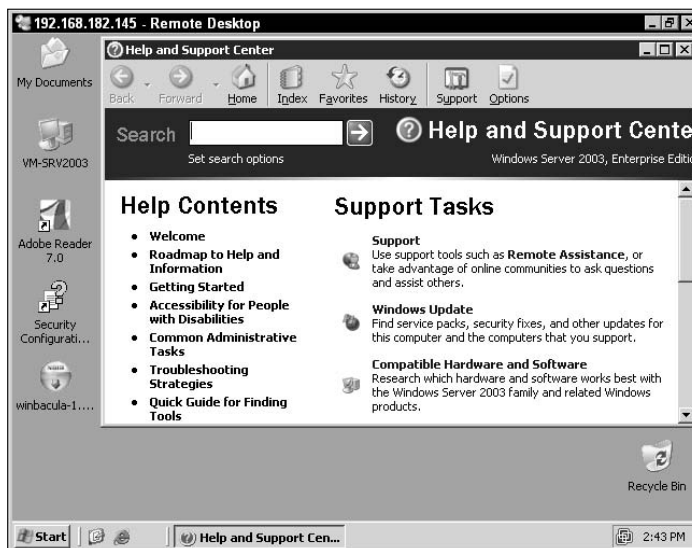


Figure 2-8:
A Remote
Desktop
Connection
to another
computer
on the
network.

Remote Desktop Connection, through RDP, is a great tool for remotely troubleshooting a variety of problems because it can give you control of a computer on a remote network as if you were sitting right in front of it.

Remote Assistance

Remote Assistance allows you to ask for or provide help to another person who is using a distant computer, by watching what they are doing, or taking control of their computer and having them watch you.

Remote Assistance uses RDP just like Remote Desktop Connection, but the implementation is very different. As a CompTIA A+ Certified Professional, you will often be asked to help a remote user perform a series of tasks. Many users just need a bit of extra guidance to complete a task, or have a question about something that is on their computer screen. After being involved in many telephone conversions with remote users, I can honestly say that a picture is worth a thousand words — sometimes even two thousand — so by having the user be able to show you exactly what is happening to them, or being able to show the user exactly what to, you can save a great deal of time.

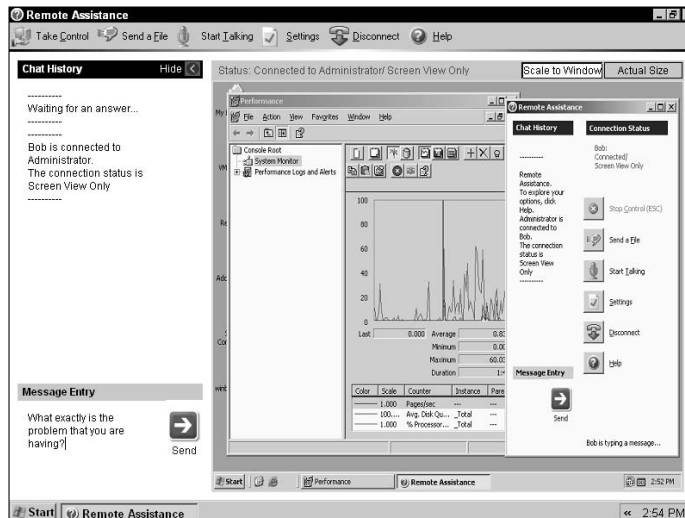
Remote Assistance starts with a user who needs help. They will go to Help and Support in their Start menu and open up the Help and Support Center (the remote help is here as well). In the Ask for Assistance section, there is an option to *Invite a friend to connect to your computer with Remote Assistance*. When they choose this and invite you to help, they will have an option of sending an invitation through Windows Messenger, e-mail, or by saving the invitation to a file. Regardless of the method they choose, an invitation file is generated, and this invitation is key to the Remote Assistance system. By opening the invitation, you will have to the ability to open the Remote Assistance application and connect to a user's computer by using RDP.



Remote Assistance requests are always initiated by the person requesting help.

The Remote Assistance application provides a section to have a chat session with the remote user, as well a section that allows you to see and/or control the remote system, as shown in Figure 2-9. Unlike Remote Desktop Connection, with this application, both the local and remote user can see what is going on with the problem computer.

Figure 2-9:
The Remote Assistance process can be a boon for many help desk personnel.



Getting an A+

This chapter takes a look at the major functions of an operating system, and in order to be prepared for the exam, you should remember the following:

- ◆ The operating system is responsible for management of the interface between all of the hardware components that make up your computer.
- ◆ General protection faults and Stop errors are often a result of applications trying to cross memory boundaries.
- ◆ Virtual memory is composed of physical memory or RAM and space on your hard drive, in the form of a paging file.
- ◆ NTFS is the most robust and stable of the available file systems but is compatible with the fewest other operating systems.
- ◆ System and application settings are stored in the Registry or in .ini files.
- ◆ `explorer.exe` makes up most graphical navigation of your computer and is represented as both My Computer and the desktop.
- ◆ Control Panel applets offer interfaces to many of the system settings for your computer.
- ◆ Remote access to computers is available through both Remote Desktop Connections and Remote Assistance, both of which rely on Remote Desktop Protocol (RDP).

Prep Test

- 1** Phil is using a Windows XP computer and has run into a problem. He has called you on the phone and has asked for you to open a Remote Assistance connection to his computer to walk him through the solution. What is the first step that you must take?
- A Open the Remote Desktop Connections program.
 - B Ask Phil to choose Help and Support from the Start menu.
 - C Ask Phil to run `geninvite.exe` from Start→Run.
 - D Open the Remote Assistance program on your computer, File→Connect to computer, and chose Phil's computer.
- 2** Virtual memory is composed of which of the following? (Choose all that apply.)
- A Physical memory
 - B Extended memory
 - C Hard drive space
 - D Compressed memory
- 3** Bill is looking to see if he has Service Pack 2 installed for Windows XP. He went to the System Information Tool located at Start→All Programs→Accessories→System Tools. Is there a quicker way he could have found out this information?
- A He could have used the Run command to run `systemtool.exe`.
 - B He could have right-clicked My Computer on the desktop or in the Start menu and chosen Properties.
 - C He could have used the command-line tool `ver` from any command prompt.
 - D There are no other locations in the OS where this information can be displayed.
- 4** You have launched an MS-DOS application on your Windows 2000 computer. When you open Task Manager to see how much processing time the application is taking, you don't see it listed. Why is that?
- A The application is running inside an NTVDM.
 - B MS-DOS applications are not listed on the Processes tab.
 - C The application is listed, but the Processes tab list is using the name stored in the header of the application rather than the executable name.
 - D You should check the list again because all running processes are listed, and you must have missed the name.

- 5** Your computer has just suffered from a Stop event (also known as the Blue Screen of Death). You are trying to decide what has caused the error. What is the most likely possibility?
- A A multimedia application
 - B A background application
 - C A network error
 - D A system device driver
- 6** Bill is setting up a computer that will be used to boot into several different operating systems. He would like to have one partition used for sharing or transferring data between the operating systems. What file system should he use?
- A FAT
 - B FAT32
 - C NTFS
 - D Reiser FS
- 7** What tool combines all of the Windows configuration and management tools into one place?
- A Computer Management
 - B Administrative Tools
 - C Microsoft Management Console
 - D Administrative Command Interface
- 8** What is the main protocol that is used by both Remote Desktop Connection and Remote Assistance?
- A Dec/Net
 - B 802.11g
 - C SNMP
 - D RDP

Answers

- 1 B. The Remote Assistance process starts with an invitation being generated by the person requesting help. This invitation is created using the Help and Support Center. *See "Remote Desktop Connection" and "Remote Assistance."*
- 2 A, C. Virtual memory is composed of both physical RAM and a swap or paging file on the hard drive. *Review "Paging your memory."*
- 3 B. You can also launch the System Information tool by typing `winmsd.exe` or `msinfo.exe` at the Run command. In addition to this, you can find the info on the System Properties, which can be opened by accessing the Properties of My Computer. If you use the `ver` command, you will only see the Windows version and not the Service Pack level. *Check out "Checking the OS version."*
- 4 A. The application is running inside one of the NTVDMs that are listed. *Peruse "Getting into the architecture."*
- 5 D. Most blue-screen errors are the result of conflicts between processes running Kernel Mode, which would include device drivers. This is not to say that the multimedia application may not have made some unsupported calls to a sound or video card to cause the problem, but this would have been the second choice. *Take a look at "Getting into the architecture."*
- 6 A. Although it has size limitations of 2 to 4GB depending on the operating system that is used, the FAT system is the most compatible with different operating systems. *Peek at "Choosing your file systems."*
- 7 C. Although Computer Management has a lot of management tools in it, it is itself a snap-in for the Microsoft Management Console (MMC). Administrative Tools are also a collection of preset MMC snap-ins. *Look over "Microsoft Management Console."*
- 8 D. The Remote Desktop Protocol (RDP) is used by both of these tools. *Study "Remote Desktop Connection" and "Remote Assistance."*

Chapter 3: Installing and Upgrading Windows

Exam Objectives

- ✓ Installing Windows 2000 and Windows XP
- ✓ Upgrading Windows
- ✓ Dual booting Windows

Although installing an operating system has gotten easier over the last few years, it is still important to understand some of the issues surrounding the installation and the upgrading of operating systems.

This chapter gives you some background on the issues that surround the installation of Windows 2000 and Windows XP. In this chapter, I walk you through these installations step by step and discuss the different phases of the installation process.

Understanding Installation Methods

You can use a number of different techniques to install an operating system. But before I discuss the actual installation procedures, you should understand the overall approach you will take to install the OS.

Boot floppy

Back in the Windows 9x days, you would install the operating system by first booting off your Windows 9x boot disk. The boot disk would contain a number of utilities to help prepare the hard drive for installation. For example, with the Windows 98 operating system, you would boot off a Windows 98 boot disk, which provides you with the `fdisk` and `format` commands so that you can partition the hard drive and then format the partitions before starting the installation. In this example, the Windows 98 boot disk also loads a CD-ROM driver so that you can access the Windows 98 CD and call the `setup.exe` program to perform the installation.

With today's operating systems, such as Windows 2000/XP/Server 2003, you can partition and format the hard drive from within the setup program when installing the operating system. This is different than Windows 9x, because

back then you had to partition the system and format the partitions first. With the newer operating system, you can perform both tasks from within the setup program. So your goal when installing newer operating systems is to simply boot from the CD and start the installation.

Today's systems can all boot from CD-ROM, but in the past, some systems could not. So the folks at Microsoft created a way to produce what is called the *setup boot disks* for Windows 2000. If you have a system that cannot boot from CD because your BIOS does not support it, you can boot off the setup boot disks that are created by `makeboot.bat`. These floppy disks will start the Windows setup, load generic drivers — including a CD-ROM driver — and then continue the setup from the Windows CD-ROM. To create the setup boot disk, you need to run the `makeboot.bat` file located in the support folder on the Windows 2000 CD-ROM from another system.

Just to stress this a little further — you need to create the Windows setup boot disk only when you have a system that does not have the ability to boot off the Windows CD-ROM.



Remember that Windows 2000 requires four diskettes when you run `makeboot.bat`. After you have created them, Windows boots off the first one and then prompts for the second, third, and fourth — be sure to label them as you create them!

Bootable CD

With the Windows 2000/XP/2003 operating systems, the CD that you use to perform the installation is a bootable CD, so you don't need to use a boot floppy. When you boot from the CD with these operating systems, the installation is activated automatically, and you instantly begin installing the operating system!

The other benefit of the Windows 2000/XP/2003 operating system installations is that Microsoft provides an opportunity to partition and format the disk from within the installation program itself! This is a huge benefit because you don't need to boot from a boot diskette first and use partitioning tools like `fdisk` as you do with Windows 9x. Your goal now is simply to start the installation and take care of all partition and formatting tasks from within the installation.

The folks at Microsoft built an extra security feature into the installation CD-ROMs for Windows 2000 and Windows XP. When the CD-ROM boots, you're asked to press a key (any key) before the installation process continues. This feature keeps you from unintentionally reinstalling the operating system if you accidentally leave the CD-ROM in the drive.

Network installation

A number of network administrators install operating systems and other applications to desktop systems via a network. When performing a network installation, you boot from a bootable floppy disk or CD, which then loads a network card driver and connects to a server on the network.

The server holds the Windows setup files needed to install an operating system to the computer. With Windows 2000 or Windows XP, these setup files are in the i386 folder located on the Windows CD-ROM. In order to perform a network installation, the i386 folder is first copied to the network server so that client computers can connect to the folder and perform the installation. After you connect to the server, you run the setup program for the operating system, and the installation is run from across the network to your system. You end up installing the OS without even touching the installation CD.



Be aware that this method of installation uses up network bandwidth and can affect the performance of the network! You also need to ensure that you have a valid license for each installation of Windows.

Drive imaging

One of today's most popular techniques for installing an OS is to use imaging software. *Imaging software* creates a snapshot of a system and stores that snapshot in an image file. When you want to configure other systems with the same setup, you apply the image (snapshot) to the other systems. The contents of the image file overwrite the contents of the hard drive of the destination systems.

Here's an example of when drive imaging should be used: Your manager asks you to install Windows XP on 20 new computers. When you receive the systems, you sit down at one system and boot off your Windows XP CD to perform a clean installation on that new computer. After the operating system is installed, you install any additional drivers and applications that users will need. For example, most users need Microsoft Office, so you install Office on the computer as well.

After you have the system properly configured, you then create an image of that system by using imaging software such as Ghost or DriveImage Pro. Typically, you would install the image on an image server so that you can connect to the image from across the network at a later time from the other systems.

When it comes time to configure the other 19 systems, you boot off a network boot disk or CD-ROM and connect to the imaging server and download

the image to each of the 19 systems. This will take far less time than performing the installation of the OS, drivers, and applications on each system individually.

Most companies today are using imaging solutions!



Again be sure that you have valid licenses for the operating system and any applications that you run. You will need to ensure that you have a license for the operating system and any applications for each system that have the image applied.

Preparing for Installation

After talking to many people in the industry about installing different operating systems or servers, I have come to realize that one thing that always requires more time is *planning*. Many people tell me that they can't do any actual work until their head office sends them the server. However, the best time to do your installation and disaster planning may be while the server is being shipped. Planning your installation or deployment of the operating system to the desktops can save you time and money in the long run by helping you to anticipate any issues that might arise and to have the solutions ready when the time comes.

This section helps you identify some points you need to consider when installing a new Windows operating system. Before jumping into the installation, you want to be sure that you plan a number of installation decisions. Some of the things you will be asked about during the installation are as follows:

- ◆ **Hardware requirements**
- ◆ **Computer name**
- ◆ **Workgroup/domain**
- ◆ **Partitions**
- ◆ **File systems**
- ◆ **Application support**

Hardware requirements

You first prepare to install an operating system by finding out the OS's hardware requirements and making sure that the computer fits the bill. Table 3-1 shows the minimum hardware requirements for Windows 2000 and Windows XP.

Table 3-1 Windows 2000/XP Minimum Hardware Requirements

<i>Hardware Component</i>	<i>Windows 2000</i>	<i>Windows XP</i>
Processor	Pentium 133 MHz	Pentium 233 MHz
Memory	64MB (128MB is recommended)	64MB (minimum) (256MB is recommended)
Hard Disk Space	650MB	1.5GB

Not only should you verify that the computer meets the minimum hardware requirements before installing an operating system, you should also make sure that the hardware components you are using will work with the operating system. For example, make sure that the make and the model number of your network card are compatible with the operating system you want to install. For years Microsoft published a list of hardware components that had been tested with each version of their operating systems. This list was known as the *hardware compatibility list*, or *HCL*. You can find the HCL for a given operating system on the Windows installation CD in a file called `hcl.txt`, or you can find the listing online at Microsoft's Web site.



It is important to note that the HCL is not a list of the only devices that will work with Windows; it is a list of devices that have been tested. If you have a network card that is not on the HCL list, it might still work with the operating system, but it hasn't been tested. There is only one way to find out whether it works — install it! You will want to make sure that you are installing it on a test system and not a production system, just in case it causes the system to crash.



The HCL has been relabeled in Windows XP — it is now known as the *Windows Catalog*. To view a list of tested products for Windows XP, check out <http://testedproducts.windowsmarketplace.com>.

Computer name

When you install Windows, you must specify a computer name for the machine you are installing. This computer name is a unique name assigned to the system and will be used to identify the system on the network. You want to plan the computer names because you are not allowed to have two Windows systems on the network that have the same name. The computer names can be up to 15 characters long, and they are not allowed to include spaces. For more information on computer names, check out the networking chapter which is Book VIII, Chapter 3.

Workgroup/domain

When installing Windows on a system, you need to specify the workgroup or domain that you want the system to be a part of. A *workgroup* is a logical

grouping of computers. For example, you could logically organize all of the accounting systems into an accounting workgroup — the benefit being that when users browse the network for resources, they can double-click the accounting workgroup to get to all of the computers in the accounting workgroup.

If you are installing Windows into a Windows domain, you will need to specify the domain name instead of the workgroup name during the installation. A *domain* is the term that Microsoft uses for its networking environments. A domain has a central server that stores all the user accounts so that users have to log on to the domain only once to access resources across the network.

If you aren't sure whether you're installing your Windows system into a workgroup or a domain, choose workgroup. You can always join the domain at a later time by going to the properties of My Computer. For more information on changing a computer name check out Book VIII, Chapter 3.

Partitions

When you install Windows, you need to partition the hard drive or select which partition you want to install Windows to. Plan this out ahead of time! *Partitions* are logical divisions of the hard drive that you can use to help separate the different types of information stored on the system. For more information on partitioning a hard drive check out Book II, Chapter 5.

When installing Windows 2000/XP, you can manage the partitions from within the Windows setup program instead of having to boot from a startup disk first and then create the partitions before installing the operating system, like you had to do in the old Windows 9x days. You can build all of your partitions from within the setup program.

Within the Windows 2000/XP setup program, you can build the partitions by choosing an area of free space and then pressing C (for create). The setup program then asks you the size of the new partition. You want to be sure that you create a partition that is large enough to hold the operating system and any future patches or updates to that operating system. For example, the laptop I am using to write this book has a 100GB hard drive. I have created a partition to hold the operating system which is 35GB in size. By the time I have installed the operating system, patched it, and installed my applications, I only have 50 percent of that space available for future updates.

The Windows setup program also offers a delete command, which lets you delete any existing partitions. After creating the partitions and choosing a partition in which to install the Windows operating system, you are asked what file system you want to use.



When you install Windows, you don't have any control over the *types* of partitions that are created during the setup. For this reason, it is best to only create one partition (to install the operating system on) during the installation, and then, after the operating system has been installed, use Disk Management to create the remaining partitions. You will get more options and flexibility from the Disk Management console than from the Windows installation program.

File systems

Windows supports three file systems: FAT, FAT32, and NTFS. The primary advantage of using the FAT file system is that it is common to all Windows operating systems. The FAT file system has its limitations, however — it can handle only up to 2GB partitions, while the FAT32 file system supports up to 2000GB partitions!



Neither version of FAT supports security features, so beware of these file systems!

NTFS (New Technology File System) is a file system supported by Windows 2000/XP/2003 that offers a richer set of features than FAT or FAT32. For example, NTFS supports setting permissions on folders and files so that you can control who can access what files. NTFS also supports auditing, compression, encryption, and quotas — a feature that allows you to limit how much hard drive space a single user can use. None of these features are available with FAT or FAT32. If you need any of the features provided by NTFS, you need to use an NTFS file system on each partition!



For the exam, be sure to be familiar with the different file systems that Windows supports. Also note that NTFS is the only file system that gives you features such as permissions, auditing, and quotas. For more information on NTFS check out Book II, Chapter 5.

Other preparation steps

Before you install Windows to your production systems you also want to make sure that you prepare for the installation by checking to ensure that the applications that are going to run on the systems work and that your CMOS is prepared to boot from CDROM. This section outlines these two additional preparation steps.

Application support

Make sure you test each application that will be running on the operating system completely to verify that it functions correctly. Back in the Windows 9x days, the operating systems were very compatible with most types of applications (there were a lot fewer software companies back then), while

Windows NT was very picky about which applications could run on it. Today, you will find that as long as you are installing a 32-bit Windows application on Windows 2000 and Windows XP, the applications will function fine. If you're installing an older application, such as an old DOS program or 16-bit Windows application, make sure that the program function correctly before installing it on a production system. The bottom line: Test everything before putting it into production! For more information on application support in Windows check out Book VI, Chapter 2.

CMOS

When you install the Windows operating system, you will boot from the Windows 2000 or Windows XP CD, and the setup program will start. Because you are booting from CD, you will need to prepare for the installation by verifying in the CMOS settings that the CD-ROM device is the first device in the startup order. This will ensure that you can boot off the Windows setup CD. For more information on configuring CMOS, check out Book II, Chapter 4.

Performing Attended Installations of Windows

This section begins with an overview of the installation process for Windows 2000 and Windows XP. The attended installation processes for these two operating systems are very similar, so if you know how to install one operating system, you will be able to install the other with no problem.

We first visit what is called attended installation of the operating systems. *Attended installations* are installations that you attend, or sit through, and answer all the questions that the setup program asks.

The Windows attended installation is broken down into three major phases:

- ◆ **Setup loader:** The setup loader phase is initiated by calling `winnt.exe` or `winnt32.exe` or by booting from the Windows 2000/XP CD-ROM. This phase copies `setupldr`, which is a mini-version of `ntldr` used by the setup program, and copies to the hard drive any files that the Windows setup utility needs.
- ◆ **Text-mode phase:** The text-mode portion of the Windows installation is typically identified by the text-based environment that has a blue background. You will recognize the text mode phase by the lack of a graphical interface. The text-mode portion of the install is controlled by a mini-kernel that is started by `usetup.exe`, which is located in the `i386` folder and called automatically.

The text-mode phase of setup is responsible for detecting basic hardware components such as CPU, motherboard, and hard drives. This phase also

creates the Registry, partitions and formats the drives, creates the file systems, and verifies that you have enough hard disk space to complete the installation. After the setup verifies that you have enough hard disk space, it copies files that are needed by the setup program to the hard disk.

- ◆ **GUI-mode phase:** After the text-mode portion of the install, the system restarts and moves into the GUI-mode phase. This phase of the setup can be quickly identified by the use of a wizard and it's Windows-like shell. The wizard asks questions; you answer each question and click Next to go to the next question.

This phase of the setup detects additional devices, installs drivers for those devices, and copies additional necessary setup files that weren't copied during the text-mode phase. During this phase, you are asked for information, such as your name, your organization, and your product key for the Windows operating system, and you are asked to agree to the end-user licensing agreement (EULA).

Performing a Windows 2000 attended installation

Before I go through how to boot from the Windows CD to start the setup program, I talk about how you can launch the Windows setup program. The installation program for Windows can be launched by booting from a CDROM or by calling the setup executables manually — `winnt.exe` or `winnt32.exe`. You execute `winnt.exe` if you are installing the operating system from a 16-bit client, such as a DOS-formatted system. If you are installing Windows 2000 on a system that has a 32-bit operating system already installed, then you will run the upgrade program, which is `winnt32.exe`. Table 3-2 and Table 3-3 show popular setup switches with these programs.

Table 3-2 winnt.exe Setup Switches

<i>Switch</i>	<i>Description</i>
<code>/U:<filename></code>	Tells the installation program the name of a file that has all of the answers to the questions that will be asked during the installation. This is used to automate the setup, also known as an unattended installation.
<code>/UDF: <id>, <filename></code>	Points to a database file (really a text file) that has a list of unique settings for each computer in the network. This is another file used in automating setup.
<code>/S:<path></code>	Specifies the source path to the Windows NT/2000 installation files.
<code>/t:: <drive></code>	Tells the setup program to store the temporary files used during the installation process on a different drive.

(continued)

Table 3-2 (continued)

<i>Switch</i>	<i>Description</i>
/r: <folderpath>	Specifies an additional folder that will be copied to the system during setup. This folder remains on the system after the installation is complete.
/r: <folderpath>	Specifies an additional folder that will be copied to the system during setup, but this folder is deleted after the setup program is completed. Using this switch is a great way to place some additional files needed by the setup program but have them removed when setup is done.
/e: <command>	Specifies a command that will execute after the graphical portion of the setup is complete.

Table 3-3 winnt32.exe Setup Switches

<i>Switch</i>	<i>Description</i>
/unattend:<filename>	Tells the installation program the name of a file that has all of the answers to the questions that will be asked during the installation. This is used to automate the setup, also known as an unattended installation.
/UDF: <id>,<filename>	Points to a database file (really a text file) that has a list of unique settings for computers on the network. This is another file used in automating setup.
/S:<path>	Specifies the source path to the Windows NT/2000 installation files.
/t:empdrive: <drive>	Tells the setup program to store the temporary files used during the installation process on a different drive.
/copydir: <folder>	Specifies an additional folder that will be copied to the system during setup. This folder remains on the system after the installation is complete.
/cmd: <command>	Specifies a command that will execute before the final phase of the setup program.
/cmdcons	Installs the Recovery Console on the system. The Recovery Console is used to help recover a system that cannot be booted. You can install the Recovery Console at any time after Windows has been installed.
/checkupgradeonly	Verifies that your existing hardware and software are compatible with the Windows 2000/XP operating system. You would use this command before you upgrade your Windows 9x system to Windows 2000 or XP.



For the A+ exam, be familiar with the setup program switches that are listed in Table 3-2 and Table 3-3.

If you want to launch the setup program by booting from the CDROM, you simply place the Windows CD in the system and power the system on. Windows will ask you to press any key if you wish to install Windows by saying “Press any key to set up Windows.” When you boot from the Windows 2000 CD, the setup program runs automatically (after you press a key), and you are installing Windows! Booting from the CD is the approach most people take when installing a single copy of Windows.



Remember that with Windows 2000/XP/2003, you don’t necessarily need to call the `winnt.exe` or `winnt32.exe` program yourself — you can boot off the CD.

If you have a bootable CD-ROM, the installation will be fairly easy. All you need to do is follow these steps:

1. Place the CD in the CD-ROM tray and boot up the system.

The setup program is invoked automatically, as shown in Figure 3-1.

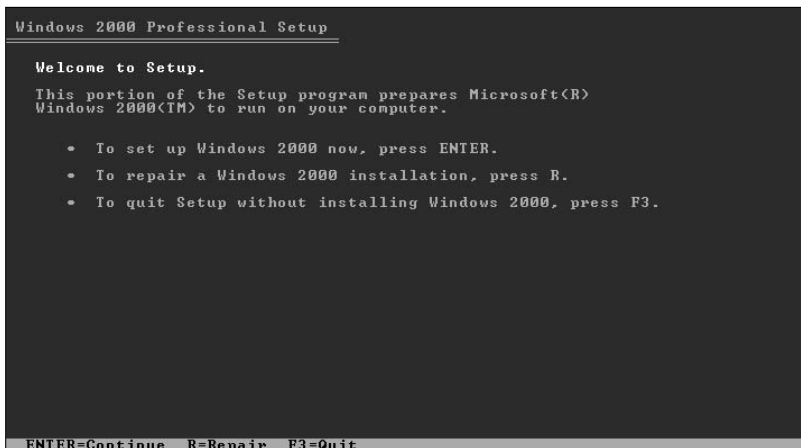


Figure 3-1:
The start of the Windows 2000 installation program.

When the setup program starts for Windows 2000, you are asked whether you want to install Windows 2000 or repair an installation of Windows 2000.

Repairing an installation of Windows 2000 is done when your existing installation has failed and you would like to fix it. In this example, though, you just want to install Windows 2000.

2. Press Enter to continue with installation.

If you are installing Windows 2000 on a system that has not yet defined any partitions, Windows 2000 will display a screen that warns you that

there is an empty disk or a disk running an operating system that Windows 2000 does not understand, as shown in Figure 3-2.

```
Windows 2000 Professional Setup

Setup has determined that your computer's startup hard disk is new
or has been erased, or that your computer is running an operating
system that is incompatible with Windows 2000.

If the hard disk is new or has been erased, or if you want to discard
its current contents, you can choose to continue Setup.

If your computer is running an operating system that is incompatible
with Windows 2000, continuing Setup may damage or destroy the existing
operating system.

• To continue Setup, press C.
  CAUTION: Any data currently on your computer's startup hard disk
  will be lost.

• To quit Setup, press F3.

C=Continue Setup  F3=Quit
```

Figure 3-2:
Partition
warning in
Windows
2000.

3. To continue with setup, press C.

You will be required to read over the license agreement and agree to its terms, as shown in Figure 3-3.

```
Windows 2000 Licensing Agreement

*****
120-Day Evaluation License for
Microsoft Windows 2000 Professional
*****

IMPORTANT-READ CAREFULLY: This Microsoft
Evaluation License Agreement ("Evaluation License") is
a legal agreement between you (either an individual or
a single entity) and Microsoft Corporation for the
Microsoft software product identified above that
accompanies this Evaluation License, which includes
computer software and may include associated media,
printed materials, and "online" or electronic
documentation ("Product"). BY INSTALLING,
COPYING, OR OTHERWISE USING THE PRODUCT,
YOU AGREE TO BE BOUND BY THE TERMS OF THIS
EVALUATION LICENSE. IF YOU DO NOT AGREE
TO THE TERMS OF THIS EVALUATION LICENSE,
DO NOT INSTALL, COPY, OR USE THE PRODUCT.

1. GRANT OF LICENSE.
Microsoft grants you the following rights provided you
comply with all terms and conditions of this Evaluation
License for the LIMITED PERIOD specified below:

a. Installation. You may install the Product on
your computers, solely for purposes of demonstration.

F8=I agree  ESC=I do not agree  PAGE DOWN=Next Page
```

Figure 3-3:
Agreeing to
the terms of
the license.

4. Read the license agreement and press F8 to agree to the terms.

You can scroll down through the agreement with the Page Down key.

After you have agreed to the licensing terms, Windows 2000 shows you a screen where you choose the partition you want to install the operating system to, as shown in Figure 3-4.

5. If you want to change your partition information by creating or deleting partitions, you can do so by pressing C (for create) or D (for delete).
6. Select the partition where you want to install Windows 2000 and then press Enter.

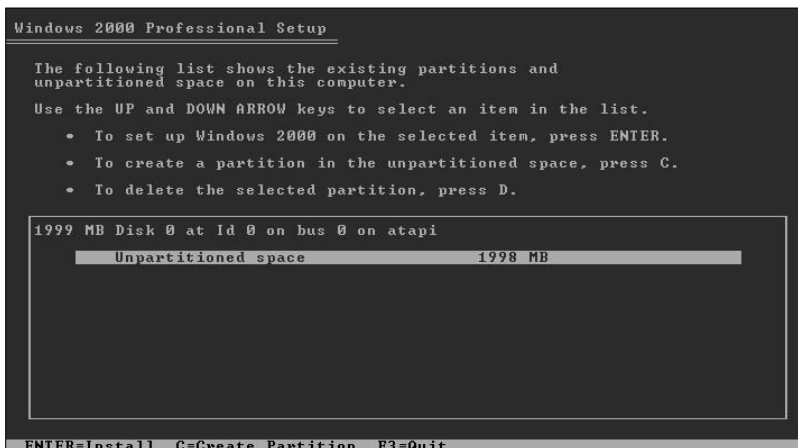


Figure 3-4: Choosing a partition to install Windows 2000 to.

After you have chosen the partition that you want to install Windows 2000 onto, you will be asked what file system you want to use for that partition, as shown in Figure 3-5.

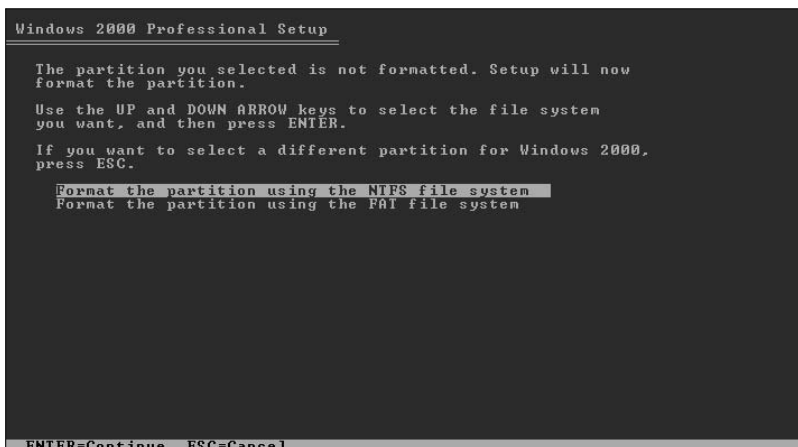


Figure 3-5: Choosing a file system during the Windows 2000 installation.



7. Select the type of file system you want to install on that partition and press Enter.

See the earlier section, “File Systems” for the pros and cons of FAT, FAT32, and NTFS file systems.

Setup will then begin copying the files that are required for the installation to your hard disk.

After the files are copied to your system, Windows 2000 will load the GUI phase of the installation before the rest of the setup will continue. During this time, the Windows 2000 setup program detects the hardware that is installed on your system. This part of the installation may take anywhere from three to five minutes, depending on your system, as shown in Figure 3-6.

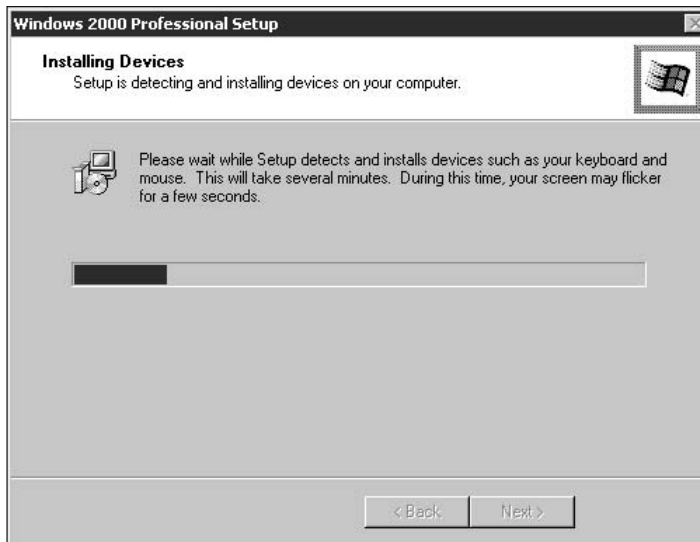


Figure 3-6:
Windows 2000 setup detects your computer hardware.

After Windows 2000 has detected the hardware in your system, it will then continue with the installation and ask for your local settings and your keyboard layout settings.

8. If you would like to change the local settings or the keyboard layout settings, click the corresponding Customize button, shown in Figure 3-7, and tweak the settings as you desire.

After you have entered the keyboard layout and the locale that you will be using, setup asks you to personalize your software. Specifically, it asks for your name and the organization name, as shown in Figure 3-8.

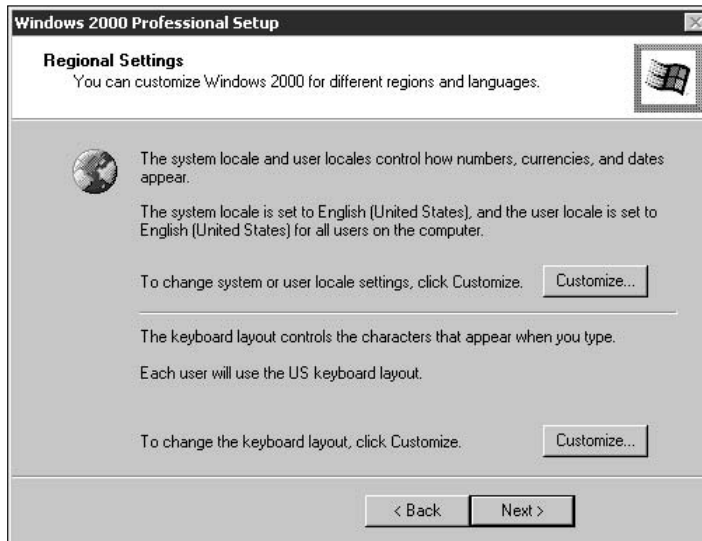


Figure 3-7: Selecting your locale and your keyboard layout.

9. Enter a name and organization name in the appropriate spaces and click Next.



It's probably a good idea to use a generic name, such as "IT Support," instead of your real name in the Name field because any new software that is installed will display this information.



Figure 3-8: Entering personalized information during Windows installation.

After you enter your personal information, setup prompts you for the product key of the operating system. You can find the product key on the CD case of Windows 2000.

10. Enter the product key and click Next.

After you enter the product key, you will be required to enter a computer name and a password for the administrator account, as shown in Figure 3-9.

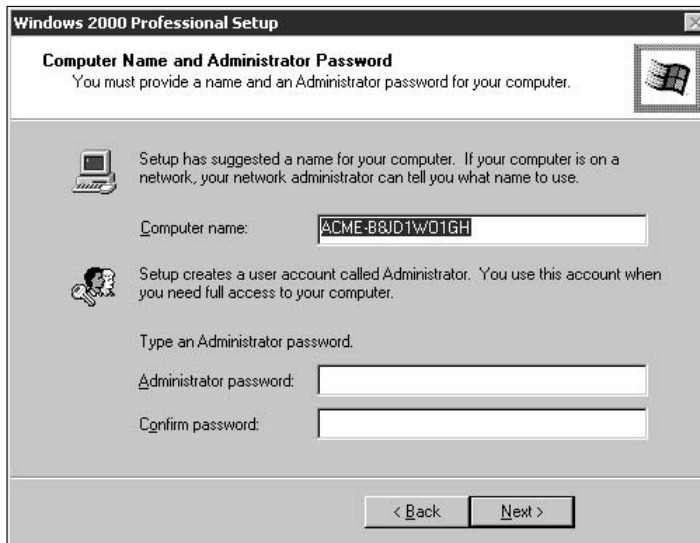


Figure 3-9: Assigning a computer name and administrator password.

11. Enter a computer name in the first text box and an administrator password in the next two boxes and click Next.

The administrator account has full access to the system. You want to create a password that is unique and difficult for someone to guess, which you can do by composing the password as a mixture of numbers and letters and also including capitalized characters at different points in the password (passwords are case-sensitive).

Note that you are not asked for the administrator's name. Each Windows installation has an administrative account that has full access to the system. When you install the operating system, that administrative account automatically has an account name `Administrator`. You are responsible only for assigning a password to that account during installation.

The next screen asks you to verify the date and time. You can also choose whether you would like the operating system to adjust the time automatically to accommodate daylight savings, as shown in Figure 3-10.

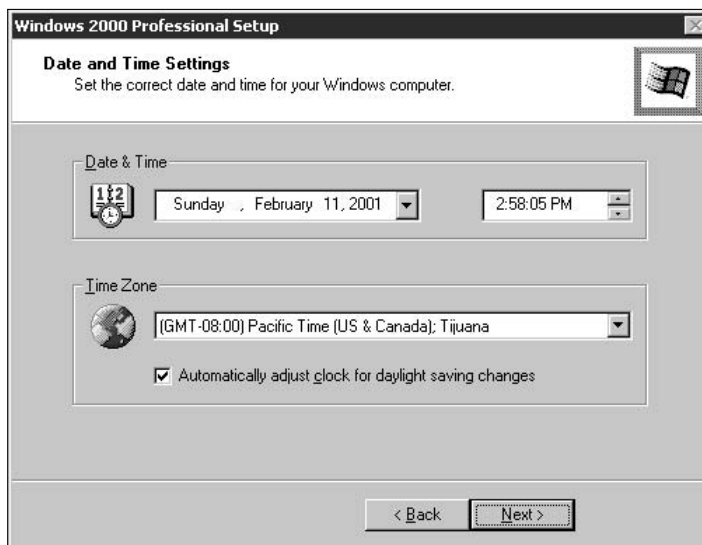


Figure 3-10: Selecting the date and time during the Windows 2000 installation.

- 12. Select the correct date, time, and time zone; also select the check box if you want the operating system to automatically adjust for daylight savings time. Then click Next.**

Next, the Installation Wizard informs you that it is configuring the networking components of the operating system. Setup then asks whether you would like to install typical network settings or to customize your network settings, as shown in Figure 3-11.

- 13. Select whether you want the typical networking settings or would like to customize the settings and then click Next.**

If you choose Typical Settings, setup installs Client for Microsoft Networks, File and Printer Sharing for Microsoft Networks, and the TCP/IP protocol configured for DHCP.

If you choose Custom Settings, you will be able to modify your networking components, such as assign yourself a static IP address. For more information on networking Windows check out Book VIII, Chapter 3.

The Installation Wizard then asks whether you want the computer to become part of a domain or whether it should be installed into a workgroup, as shown in Figure 3-12.

Figure 3-11: Choosing a network installation type.

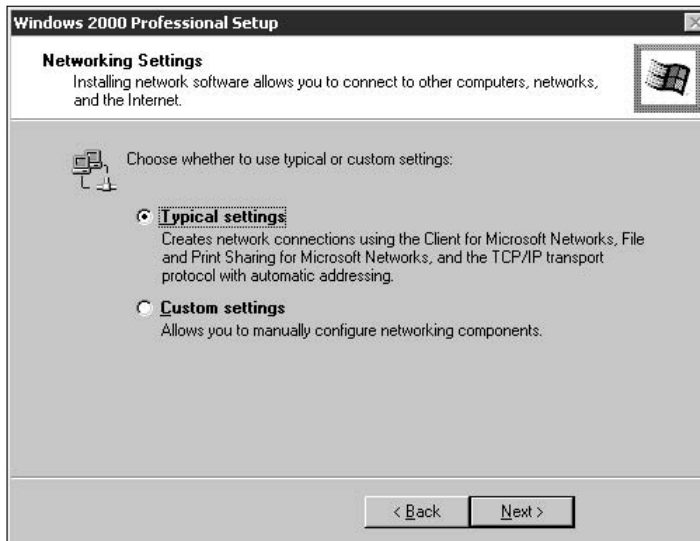
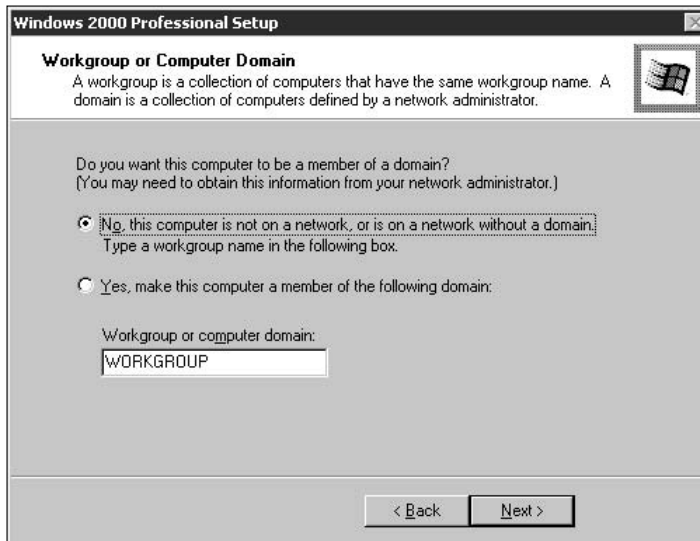


Figure 3-12: Selecting your workgroup or domain during installation.



14. If your computer will be in a workgroup, select the first radio button; if it will be in a domain, select the second one. Then enter the workgroup or domain name and click Next.



The difference between a domain and a workgroup is that the domain has a central database of users and computers for people who wish to log on to the network and use certain network resources, such as a printer. A workgroup environment doesn't have a server with a central list of user accounts — each workstation manages its own user accounts in the local SAM database.

If you aren't sure whether your computer is in a domain environment, choose workgroup. You can always join the domain later on.

Windows continues by installing all of the Windows components that are required by the operating system for day-to-day use. It then sets up the Start menu, registers components on the machine, and finally cleans out any temporary files that were created by the installation program, as shown in Figure 3-13.

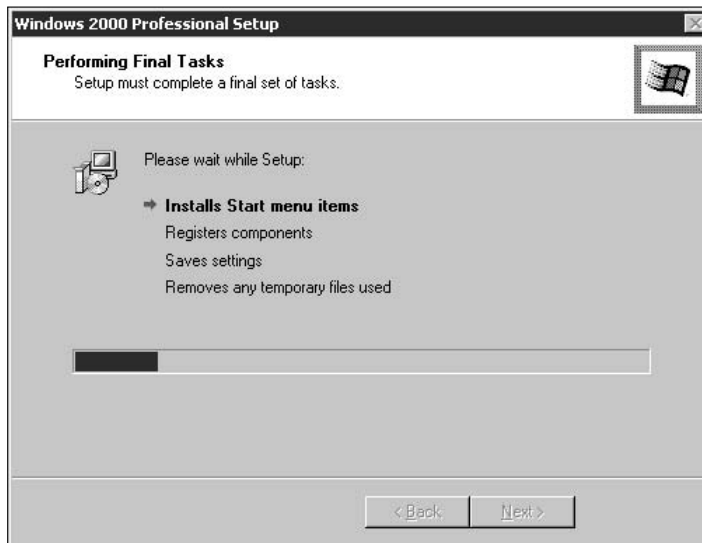


Figure 3-13:
Finalizing
the
Windows
2000
installation.

15. After the final stage has completed, click the Finish button to restart the computer.

Upon restart, Windows 2000 takes you into the Network Identification Wizard, where you tell Windows 2000 whether you will be using the same username and password to log on each time or whether each user should provide a unique username and password.

16. Select the appropriate logon option and click Next.

If you tell Windows to log on with the same account each time, Windows will automatically log on — meaning whoever sits at the computer and turns it on isn't prompted for a username and password. If you would rather the system require you to log on manually each time, you should select the Users Must Enter a User Name and Password to Use This Computer option, shown in Figure 3-14.

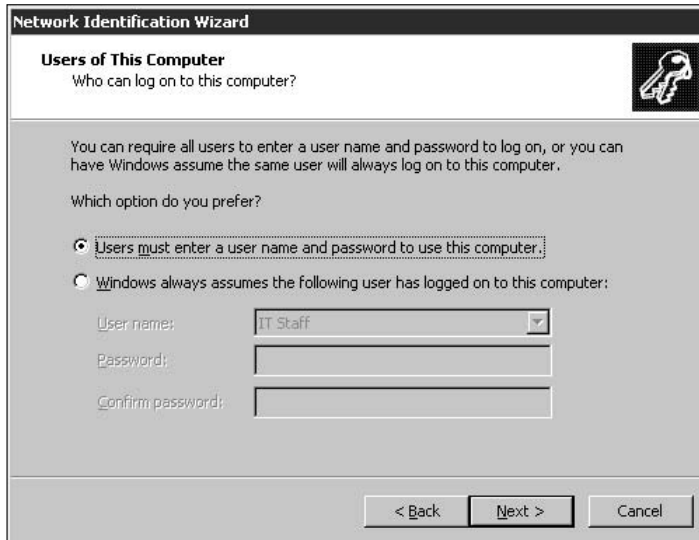


Figure 3-14:
Configuring
Windows
2000 for
logon type.

After you finish the Network Identification Wizard, the installation of Windows 2000 is complete. Congratulations!



To practice installing Windows 2000, check out Lab 3-1. Lab 3-1 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Performing a Windows XP attended installation

Installing Windows XP is similar to installing Windows 2000 because both operating systems are based on the old Windows NT technologies. This means that you will see a lot of similarities between the two operating systems — starting with the installation.

To perform a Windows XP attended installation, follow these steps:

- 1. Put the Windows XP installation CD in the CD-ROM drive and restart the computer.**

2. When the screen says, “Press any key to boot from CD,” press a key.

After you boot from CD, the *text-mode* portion of the installation begins. You can always tell when you’re in the text-mode portion of the installation because of the blue background and the lack of the “Windows-ish” look.

After booting from the CD-ROM, setup copies temporary files to the hard disk and then asks you whether you want to install Windows or recover Windows, as shown in Figure 3-15.




Figure 3-15:
Just press Enter to start the installation.

3. Press Enter to start the Windows XP installation process.

4. Read the licensing agreement (shown in Figure 3-16) and then press F8 to agree to the terms.

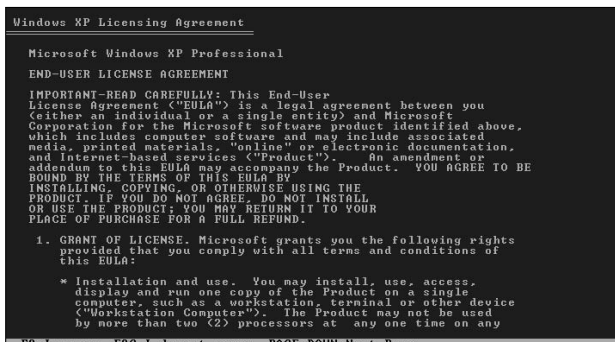



Figure 3-16:
Agreeing to the licensing terms.

You are then asked to partition the disk. Like with the Windows 2000 installation, you can manage partitions here and perform actions such as deleting and creating partitions.

5. To create a partition, press C, type 8000 (MB) as the partition size, as shown in Figure 3-17, and then press Enter.


Figure 3-17:
Creating an 8000MB partition to install Windows XP.



6. Select the newly created partition to install Windows XP to and then press Enter.

You are asked what file system you would like to format the partition with. You may format the partition with the FAT file system or NTFS, as shown in Figure 3-18.

Figure 3-18:
Performing a quick format and applying NTFS to the partition.



7. Choose to format for NTFS (Quick) and press Enter.

The difference between a quick format and a regular format is that with a quick format Windows does not perform a scan for bad sectors, whereas a regular format does do the scan for bad sectors. Scanning for bad sectors can be time-consuming, so we have chosen to perform a quick format.

Setup continues by formatting the partition and then copies files to the Windows folder on the hard disk. After the files are copied to the hard disk, setup reboots and the GUI-mode portion of the installation starts, as shown in Figure 3-19.



Figure 3-19:
The GUI-mode setup phase of the Windows XP installation.

Setup installs devices that are detected during installation, and then asks you to choose your Regional options, such as language settings.

8. Choose Next to accept the default regional settings.

We will accept the default settings in order to complete the installation. You may change your regional settings at any time through the Windows Control Panel.

9. Fill in “IT Support” as your name and then type your organization name, as shown in Figure 3-20, and then click Next.

10. Type the product key (found on the CD case that the installation CD came in) and then click Next.

11. Enter a computer name and type (twice) an Administrator password; then click Next.

The administrator account has full access to the system, so you want to supply a strong password that incorporates letters (both upper- and lowercase) and numbers.

12. Choose your time zone and then click Next.

You are then asked about network settings.



Figure 3-20: Supply your name and organization to the Windows XP setup program.



13. Click Next to accept the typical network settings.

If you install the typical network settings, you will have File and Printer Sharing, TCP/IP, and the Client for Microsoft network installed.

14. Click Next to accept the default workgroup for Windows XP.

Windows continues with installation and finally reboots. After the reboot, Windows finalizes the setup by walking you through a configuration wizard, as shown in Figure 3-21.

Figure 3-21: Finalizing the Windows configuration after Windows has been installed.



The wizard first asks whether you want to connect to the Internet through the LAN setup.

15. If your computer will connect to the Internet through the network, choose this option and click Next.

You may then choose whether you wish to register Windows with Microsoft.

16. Choose your registration preference and click Next.

17. Type the name of a username to use as your own user account and then press Next to finalize the setup.

Because you have not been given the opportunity to assign a password for this user account, you will log on without a password for this account until you assign the password through the User Accounts applet in the Control Panel.

When you install Windows XP, you are also asked to prove the validity of your Windows installation by *Activating* Windows. When you activate your installation of Windows, no other installation of Windows can be activated with that copy of Windows. You may choose to activate Windows via the Internet or via telephone by calling the phone number provided in the wizard.

Performing Unattended Installations of Windows

The opposite of an attended installation is an unattended installation. Doing an *unattended* installation means that you don't have to be present when the installation is run because you place all the answers to the installation questions in a text file, and the setup program reads the text file.

When you need to install the operating system on many computers, you can do an automated, *unattended* installation. You start this process by creating an *answer file*, also known as an *unattend file*, that has all of the answers to the questions you are asked during the installation. The Windows 2000/XP/2003 setup program then uses the answers in this file, instead of prompting you for each answer, during installation. A typical answer file is shown in Figure 3-22.

The answer file is divided into different sections, which are indicated by square brackets. For example, notice in the following sample that there is a section called [UserData]. In this section, you can set things like what you want the setup program to use for your name and organization and the computer name of the system being installed.

```
[UserData]
  ComputerName=WORKSTATION1
  FullName="Glen E. Clarke"
  OrgName="CompanyABC"
```

Figure 3-22:
Looking
at the
structure of
an answer
file.



```
companyabc.txt - Notepad
File Edit Format View Help
:SetupMgrTag
[Data]
  AutoPartition=1
  MsDosInitiated="0"
  UnattendedInstall="yes"

[Unattended]
  UnattendMode=FullUnattended
  OemSkipEula=yes
  OemPreInstall=NO
  TargetPath=%WINDIR%

[GuiUnattended]
  AdminPassword="password"
  EncryptedAdminPassword=NO
  OEMSkipRegional=1
  TimeZone=50
  oemSkipWelcome=1

[UserData]
  ProductID=????-????-????-????-????
  FullName="Glen E. Clarke"
  OrgName="CompanyABC"
  ComputerName=WORKSTATION1

[TapiLocation]
  AreaCode=902

[Branding]
  BrandIEUsingUnattended=yes

[URL]
  Home_Page=http://www.glenecClarke.com
```

In the preceding sample, `ComputerName` is a setting that corresponds to a question asked during the installation, while the label `WORKSTATION1` is the value you are assigning to that entry.



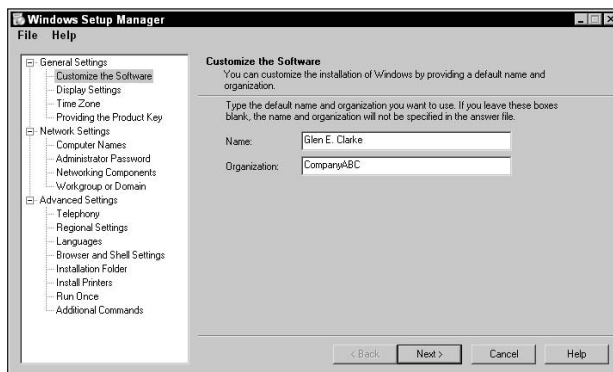
To get a list of settings that can be assigned in the answer file, look at the help file located in the `deploy.cab` file, which you can find on the Windows CD-ROM in the `Support\Tools` folder.

To make it easier to create the answer file, Microsoft has created a wizard, known as the *Setup Manager* (`setupmgr.exe`), that you can run. You can locate the Setup Manager in the `deploy.cab` file on the Windows CD as well.

After extracting the Setup Manager from the `.cab` file, you can run it by double-clicking `setupmgr.exe`. The wizard launches, asks all of the questions you would be asked during a typical installation, and then creates the answer file. Figure 3-23 shows the Setup Manager asking for your name and organization. Notice that after you answer a question, you click the **Next** button to get the next question — a typical wizard format.

After you have created the answer file, you will then pass it to the Windows setup program by calling the `winnt.exe` or `winnt32.exe` program and supplying the switch that tells the program what answer file you would like to use. Do you remember what switch to use to specify an unattended file for `winnt.exe`? If you answered `/u`, then you are correct!

Figure 3-23:
Running
the Setup
Manager for
Windows
simplifies
creating
answer
files.



For the exam, remember that setup manager is a Windows program that is used to create answer files. Once the answer file is created, you can then pass the answer file to the Windows setup program by using `winnt.exe /u: <filename>` or `winnt32.exe /unattend: <filename>`.

Assume that you've copied the CD contents to a folder on the server and have created a network drive (J: in this example) to point to that folder. To perform an automated installation of Windows, connect to the J:\i386 directory and upload your newly created answer file (`companyabc.txt` for the sake of this example). To start the installation process from a DOS-formatted system that needs Windows 2000 installed, type the following at the command prompt:

```
J:\i386\winnt.exe /s:j:\i386 /u:companyabc.txt
```

The preceding example code will run the `winnt.exe` file and supply J:\i386 as the source directory (where the Windows files are located) to the setup program. The `/u` switch tells the `winnt.exe` program to use the `companyabc.txt` file as the answer file — pretty cool technology!



To install Windows to a system running a 32-bit operating system, you run the `winnt32.exe` program. Also note that the switches are a little bit different as well. For example, you use the `/unattend` switch instead of the `/u` switch.

If you don't have a network, and you want to automate a Windows installation from CD, you can do that! You simply create an answer file (on a computer that is already running), name it `winnt.sif`, and copy that file to a floppy disk. When you boot from CD to install Windows, pop the floppy disk in really quick — the setup program will automatically read the `winnt.sif` file, and you are off to the races with an automated, unattended installation!



To practice automating a Windows XP installation, check out Lab 3-2. Lab 3-2 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Verifying and Troubleshooting the Installation

In the previous two sections, I show you how to perform an attended and unattended installation of Windows. In this section, you look at some tips for verifying and troubleshooting a Windows installation.

After Windows has been installed, the first thing you want to do is ensure that all device drivers have been installed. To do this, go to Device Manager (choose `Start` → `Control Panel` → `System` and then click the Device Manager button on the Hardware tab) and verify that all drivers are present. If you are missing drivers for some devices, the Device Manager will have a category called Unknown Devices, which are indicated by a yellow question mark.



Check out Book VI, Chapter 1 to find out more about updating drivers.

Windows creates a number of log files during installation. If installation doesn't go as you planned, you can read these logs to get more information about why an error occurred. Keep in mind that these log files report information only for events that happen during the installation, so if you're troubleshooting something like booting Windows or networking, these files won't help.

The log files created during the installation are:

- ◆ `Setupact.log` records information about the files that are copied during the installation.
- ◆ `Setuperr.log` records information about errors that happen during the installation.
- ◆ `Setupapi.log` records information about device driver files that are copied during the installation.
- ◆ `Setuplog.txt` records additional information about the device drivers.



Be sure to remember the log files that are created when you install Windows. These files can be used to help troubleshoot problems with the installation of Windows.

A number of problems may arise during Windows installation. The following list outlines a few common ones:

- ◆ **Bad CD-ROM:** One problem might be the failure to copy files during the setup process. If this happens, you might have a bad CD-ROM device, or

maybe the CD itself is bad. Try booting from another Windows CD to find out whether the problem is with the media or the CD-ROM device.

- ◆ **Can't boot after install:** If you are having trouble rebooting after you have completed the installation, be sure that you haven't left a floppy diskette in the floppy drive. Also check CMOS to ensure that the hard drive appears as a boot device. (See Book II, Chapter 4, to find out how to tweak CMOS settings.)
- ◆ **Can't start Windows install:** If you can't start the Windows setup by booting off the CD-ROM, make sure that the CD-ROM device is listed as a startup device in CMOS. Also, ensure that the CD-ROM is listed before the hard disk.
- ◆ **Not enough disk space:** The Windows setup program checks available hard disk space. If your computer doesn't have enough space, Windows indicates this with an error, and the Windows installation stops. You will need to restart the installation again after cleaning up disk space.

Overall, Windows installations have become more reliable with each successive version, so you are less likely to have installation problems with the current version than you would have back in the Windows 9x days.

Upgrading Windows



You will be responsible for understanding how to upgrade from previous versions of Windows to Windows 2000/XP. The following sections introduce you to the theory of upgrading an operating system.

Preparing to upgrade

Many of today's businesses run a Windows operating system on their computers, so it's important to understand how to upgrade from a previous Windows operating system to a more current one. Whether you are installing the operating system from scratch or performing an upgrade, the installation process is pretty much the same, although, with an upgrade, you'll find that many installation decisions were already made for you during the installation of the previous OS. For example, when you upgrade a Windows NT 4.0 computer to Windows 2000, you won't be asked the name of the computer to be installed — this information is inherited from the previous operating system.



Before you upgrade the operating system, back up *the system and any important data*. You never know when an upgrade is going to go bad, so make sure you can at least put the system back to its previous state by performing a backup before you attempt to do the upgrade.

Upgrading to Windows 2000 Professional

To upgrade to Windows 2000 from a previous version of Windows, run the Windows 2000 upgrade program (`winnt32.exe`). If you would like to check the current operating system for known problems with the installed software and hardware, you may run the upgrade program with the `/checkupgradeonly` switch. This switch launches the Windows 2000 Readiness Analyzer, which will report any hardware or software compatibility problems Windows 2000 might have with your current system. With the readiness analyzer, you can save the report to file so that you can print it and deal with the installation issues.

Generally, you will have no problem upgrading Windows 95 or Windows 98 to Windows 2000, but due to the security changes in the Windows 2000 operating system, some of the software that ran fine on Windows 95 and Windows 98 won't run on Windows 2000. If you find that this is the case, you could try to find an equivalent application that will run on Windows 2000 and install it after you upgrade the OS.



For the A+ exams, it is important to know the upgrade path from each operating system to the Windows 2000 operating system, as shown in Table 3-4. Windows 95/98/NT can be upgraded directly to Windows 2000, while Windows 3.1 needs to be first upgraded to Windows 95 or Windows 98 and then upgraded to Windows 2000 Professional. Because Windows ME came out after Windows 2000, there is no upgrade path from Windows ME to Windows 2000.

Table 3-4 Client Upgrade Paths to Windows 2000 Professional

<i>Current Operating System</i>	<i>Upgrade To</i>
Windows 3.1	Windows 95/98 and then to Windows 2000 Professional
Windows 95	Windows 2000 Professional
Windows 98	Windows 2000 Professional
Windows NT Workstation 4.0	Windows 2000 Professional
Windows ME	Cannot be upgraded to Windows 2000



To find out whether there's any incompatible software or hardware before upgrading to Windows 2000, you can run the setup in the check upgrade mode by typing `winnt32.exe /checkupgradeonly`.

When you upgrade from Windows NT 4.0 to Windows 2000, you likely won't run into any major upgrade issues because these versions have a similar structure. The similarities between Windows NT 4.0 and Windows 2000 include the following:

- ◆ Registry structure
- ◆ File systems
- ◆ Security

To upgrade from Windows NT 4.0 to Windows 2000, you run the `winnt32.exe` program. `winnt32.exe` is the upgrade program designed to run from a 32-bit operating system for the purpose of upgrading to Windows 2000. When you run `winnt32.exe`, many of the operating system's settings are carried into the installation program from Windows NT 4.0 — this means that you will have fewer questions to answer during the installation.



Be sure to review the hardware requirements of Windows 2000 before you upgrade and make sure to set up a test environment in which all applications and existing hardware have been thoroughly tested.

Upgrading to Windows XP

A number of previous operating systems can be upgraded to Windows XP, including Windows 98, Windows 98 SE, Windows ME, Windows NT 4.0 (SP5 or higher), and Windows 2000. You can't upgrade Windows 95 or Windows NT 3.51 to Windows XP directly — you must first upgrade those operating systems to Windows 98 (for existing Windows 95 systems) or Windows NT 4.0 SP5 (for Windows NT 3.51). Table 3-5 gives a summary of the operating systems that can be upgraded to Windows XP.

Table 3-5 Client Upgrade Paths to Windows XP

<i>Current Operating System</i>	<i>Upgrade Directly to Windows XP?</i>
Windows 3.1	No
Windows 95	No
Windows 98	Yes
Windows NT Workstation 4.0	Yes (Windows NT 4.0 with service pack 5)
Windows ME	Yes
Windows 2000	Yes

Before you upgrade Windows XP, you can check the existing system for known compatibility issues with Windows XP by using the Upgrade Advisor. The Upgrade Advisor identifies compatibility issues and gives details on resolution for these issues.

To run the Upgrade Advisor, execute `winnt32.exe` with the `/checkupgradonly` switch, or you can launch it from the splash page that appears after inserting the Windows XP CD. After loading the Windows XP

CD-ROM, choose the Check System Compatibility option and then choose to Check My System Automatically. The Upgrade Advisor will run and will display any compatibility problems, as shown in Figure 3-24.

Figure 3-24:
The Upgrade Advisor identifies known hardware and software issues when upgrading to Windows XP.



Notice in Figure 3-24 that the Upgrade Advisor has identified an issue upgrading the system because the existing system has CD Creator installed. You can save the report to a file by pressing the Save As button, or you can view the details in an HTML page by choosing the details button, as shown in Figure 3-25.

Figure 3-25:
Displaying the details of a report created by the Upgrade Advisor.



Installing Additional Windows Components

After installing Windows 2000 or Windows XP, you can add or remove components (or portions of the Windows operating system) that have been

installed by going through Add/Remove Windows Components. For example, if you decide you would like to create a Web site, you will most likely want to have a Web server installed on your system. Windows 2000 Professional and Windows XP don't come with Web server software installed, so if you want it, you will need to install it after you install the operating system.

To install additional Windows components after the installation of the operating system, follow these steps:

1. Choose Start→Settings→Control Panel (in Windows 2000) or Start→Control Panel (in Windows XP).

2. In the Control Panel, click Add or Remove Programs.

In the Add or Remove Programs dialog box, you will notice on the left the option to Add/Remove Windows Components. This option lets you add software that comes with Windows but that wasn't installed by default, such as accessories (including games) and Internet Information Server (Microsoft's Web server software).

3. Click Add/Remove Windows Components.

The Add or Remove Windows components dialog box appears as shown in Figure 3-26.

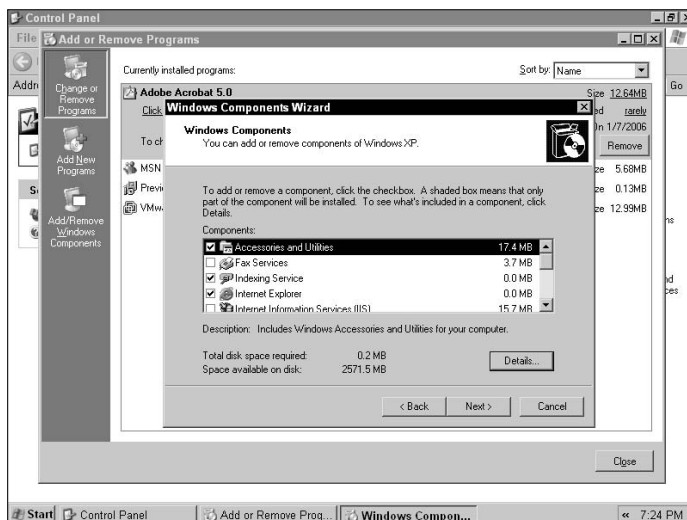


Figure 3-26:
Installing additional Windows components.

4. Select a component to install by checking the checkbox for the component and click Next.

The files are copied to your system from the Windows CD, so be sure to have the Windows CD close by in case you are prompted to insert the CD.

5. After the files have been copied, click Finish.

6. Close the Add/Remove Programs dialog box.

Understanding How to Dual Boot Windows

Dual booting is the concept of running multiple operating systems on the same computer. Dual booting is different than performing an upgrade — with an upgrade, all of your applications and their settings carry forward into the new operating system. When you dual boot, you install each operating system into a different folder, which means that the applications don't carry forward into your new operating system. With a dual boot scenario, you will need to install the applications with each operating system that is installed.

One reason you may want to dual boot multiple operating systems is to test or support applications in the different operating systems. You may also want to dual boot if your company uses an application that won't function on your primary operating system (for example, Windows 2000) but works great in an older operating system (for example, Windows 98). You can install both operating systems on the computer and then install the application that doesn't work in Windows 2000 on the Windows 98 system; any time you want to use that application, you just boot to Windows 98.

In order to dual boot multiple operating systems, the following criteria must be met:

- ◆ **The bootable drive (usually drive C) must have a file system supported by all operating systems on the computer.**
- ◆ **You must install each operating system into its own folder.**

Here's an example: If you would like to dual boot Windows 98 and Windows 2000 on the same computer, the first thing you need to do is make sure that your bootable partition (drive C) is using the file system common to all operating systems (in this example, FAT32). After you have ensured that you are using a common file system, you may then install each operating system. Remember that you will install each operating system into a different directory. For this example, assume you have installed Windows 98 in C:\windows. You can then install Windows 2000 to, for instance, C:\win2k.

You now have a system that dual boots between Windows 98 and Windows 2000. So how do you select which operating system you want to boot to? The Windows 2000 boot menu displays the different operating systems; when it appears, you select one of the operating systems.



In dual boot scenarios, it is important to note that any applications you want to use in both environments must be installed on both operating systems. For example, if you would like to use Microsoft Word in both Windows 98 and Windows 2000, you need to install it in both operating systems.

Updating Installation Files

A number of organizations copy the contents of a Windows CD to a folder on a server and then install Windows from that folder by calling the `winnt.exe` or the `winnt32.exe` setup programs. This folder is called a *distribution point* because it is used to distribute the operating system out to the client systems.

Eventually, Microsoft will release Service Packs that you will apply to the operating systems on each computer, but you will also want to update the distribution point so that it contains the Service Pack updates for any new systems that install from the folder. The benefit is that, after the folder has been updated with the Service Pack, when new systems install the operating system over the network, it will already include the updated Service Packs. Updating the source files of an operating system at a distribution point to include Service Pack files is known as *slipstreaming*.



Slipstreaming is the term used when you update your Windows source files with service pack files. The benefit is that any new installation of Windows from the updated source files will already have the service pack installed.

To slipstream a distribution point, follow these steps:

- 1. Create a distribution point on a server by creating a folder.**

For example, assume that the `K:` drive is the network drive that will contain the Windows 2000 source files. You might create a folder called `K:\win2000\i386`.

- 2. Copy the contents of the i386 folder from the Windows 2000 CD to**

`K:\win2000\i386`.

- 3. Download the proper Service Pack.**

Assume that you want to update the `i386` folder with the Windows 2000 Service Pack 4 files. You would download the executable for Service Pack 4 (`w2ksp4.exe`).

- 4. Run the service pack's executable and pass it the `/s` switch to indicate the source folder to update, like this:**

```
w2ksp4.exe /s:K:\win2000\i386
```


After the source files on your distribution server have been updated, anyone who installs Windows from that folder by calling the `winnt.exe` or `winnt32.exe` command will automatically have Service Pack 4 installed with the operating system!

Restoring User Data Files

You can move a user's computer's state from one system to another. The *state* of a user's computer is all of the information and settings that are important to that user. This includes files in the My Documents folder, e-mail settings, Internet Explorer favorites, and the desktop wallpaper, just to name a few.

Saving a computer's state

If you're going to replace a user's computer, it's important that you know how to save the computer's state. You may be replacing the system for a number of reasons. Maybe the system is running Windows 98 and you feel that, to run Windows XP, the user should have up-to-date hardware. Or maybe the user is running Windows XP on a system and your manager has a new laptop for the user that will run Windows XP. Bottom line, you have to move the settings from one computer to another.

Windows XP has a feature called the User State Migration Tool (USMT). The *USMT* is a set of features that can simplify your life when you need to move computer settings from one system to another. The USMT is made up of the Files and Settings Transfer (FAST) Wizard and the command-line tools.

To save a computer's state, follow these instructions:

- 1. Activate the Files and Settings Transfer Wizard by choosing Start⇨ All Programs⇨ Accessories⇨ System Tools⇨ Files and Settings Transfer Wizard.**

When you run the wizard, it asks you whether you're on the old computer or on the new computer (see Figure 3-27).

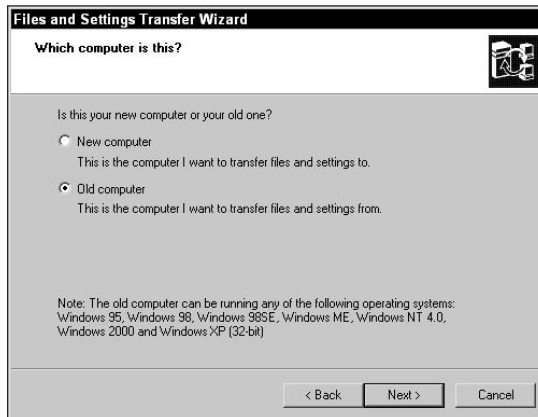
- 2. Select the Old Computer option and click Next.**

After you let the FAST Wizard know that you're on the old computer, it asks how you want to transfer the settings.

- 3. Specify where the wizard should save the computer's files and settings and click Next.**

You can specify a floppy disk or a folder in which to store the settings. If you choose a folder, you can specify to write to a folder on either a removable drive or a network drive.

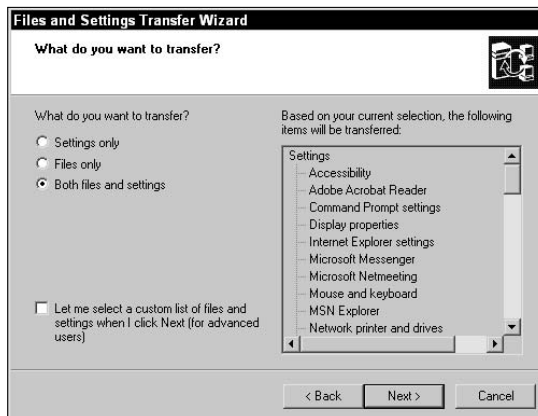
Figure 3-27:
Files and
Settings
Transfer
(FAST)
Wizard
copies a
user's
settings
from one
computer to
another.



4. Select the appropriate option for saving only the user's settings, only the user's files, or both and then click Next.

See Figure 3-28. Normally, you'd choose Both Files and Settings.

Figure 3-28:
Choosing to
save both
the files and
settings
within the
FAST
Wizard.



The FAST Wizard copies all the user state information that you specified on that computer. If you selected the Both Files and Settings option, the FAST Wizard saves these items by default to the location you specified:

- Contents of My Documents
- Contents of My Pictures
- Contents of Desktop
- Contents of Favorites

- Internet Explorer Favorites
- Browser and mail settings
- Accessibility options
- Display properties
- Folder and Taskbar settings
- Mouse and Keyboard settings
- Regional settings
- Microsoft Office applications settings

5. After the FAST Wizard saves the computer's state, click Finish to exit the wizard.



The Files and Settings Transfer Wizard allows you to migrate only one user's settings at a time. To migrate a number of user's settings simultaneously, use the `scanstate` and `loadstate` executables. `scanstate` captures a user's settings from a source computer and dumps the settings to an intermediate folder. `loadstate` loads those settings from the intermediate folder to the destination system.

Restoring a computer's state

When it comes time to restore the settings to the destination system, if the destination system is running Windows XP, simply start up the FAST Wizard again (Start⇨All Programs⇨Accessories⇨System Tools⇨Files and Settings Transfer Wizard) and select the New Computer option. Just provide the location of the user's state, and the state will be restored.



When replacing an older system, such as a Windows 98 system, with Windows XP, you will need to back up the user settings and files. You can use the Windows XP Files and Settings Transfer Wizard to do this! You will need to take the Windows XP CD-ROM to the Windows 98 system and place it in the CDROM tray. After loading the Windows XP CD in the CDROM tray, `autorun` should kick in and display a list of options on a splash page that you can perform with the CD. To run the FAST wizard you choose *Perform Additional Tasks* and then choose *Transfer Files and Settings* from the splash screen displayed when you insert the CD.

Getting an A+

This chapter introduces you to some of the guidelines for installing the different Windows operating systems. Some of the key points to remember are:

- ◆ When planning your installation, be sure to plan the partition size, file system, and the computer name that you will use when installing Windows.
- ◆ You can boot off the Windows 2000/XP CD-ROM to invoke an installation.
- ◆ You should launch the `winnt.exe` setup program from a DOS system; launch `winnt32.exe` from an existing 32-bit Windows interface such as Windows 9x.
- ◆ To dual boot multiple operating systems, make sure drive C uses a file system that is understood by each operating system and then install each operating system into a different folder.
- ◆ You can slipstream your distribution folder so that new installations will deploy with the current Service Packs already installed.
- ◆ The user state migration tools are used to copy a user's settings and files from one computer to another. A great feature to make your life easier if you need to do this some day!

Prep Test

- 1** What command starts a Windows 2000 installation from a DOS prompt?
 - A server.exe
 - B winnt.exe
 - C workstation.exe
 - D setup.exe

- 2** What term describes updating the Windows source files on a server to contain Service Pack updates?
 - A service packing
 - B Windows Update
 - C slipstreaming
 - D patching

- 3** What setup switch verifies that existing hardware and software will work with Windows 2000/XP?
 - A /checkcompatibility
 - B /verifyhardware
 - C /verifycompatibility
 - D /checkupgradeonly

- 4** What setup switch points the Windows installation to an answer file?
 - A /U: <file>
 - B /UDF:<file>
 - C /UDF:number,<file>
 - D /A:<file>

- 5** What is the minimum processor speed for Windows 2000?
 - A 233 MHz
 - B 133 MHz
 - C 266 MHz
 - D 100 MHz

- 6** Where would you go to install additional operating system components after the installation is complete?
- A Add/Remove Hardware.
 - B System icon.
 - C Add/Remove Programs.
 - D You can't install additional operating system components after the OS is installed.
- 7** What are the minimum RAM requirements for Windows 2000?
- A 32MB
 - B 64MB
 - C 128MB
 - D 168MB
- 8** What switch on the Windows setup program allows you to pass a file that contains unique settings per computer?
- A /u
 - B /upgradeonly
 - C /udf
 - D /a
- 9** After installing Windows XP, what will you need to do to prove you have a valid copy of XP?
- A Register it
 - B Restart it
 - C Activate it
 - D Shut it down
- 10** What are the recommended RAM requirements for Windows XP?
- A 32MB
 - B 64MB
 - C 128MB
 - D 256MB

Answers

- 1 B. The command to start a Windows 2000 installation is `winnt.exe`. See *"Performing a Windows 2000 attended installation."*
- 2 C. When you want to update the source files on the server, you can run the service pack executable with a `/s` switch — this is known as *slipstreaming*. Review *"Updating Installation Files."*
- 3 D. The `/checkupgradeonly` switch is used by the `winnt32.exe` program to verify compatibility with existing hardware and software. Check out *"Upgrading Windows."*
- 4 A. To point Windows to an answer file, run `winnt.exe` with a `/U:` switch followed by the path to the file that has the answers to use for the installation. Peruse *"Performing Unattended Installations of Windows."*
- 5 B. The minimum processor requirement to install Windows 2000 is a Pentium 133 MHz processor. Take a look at *"Performing Attended Installations of Windows."*
- 6 C. To install additional Windows components, go to Add/Remove Programs in the Control Panel. Peek at *"Installing Additional Windows Components."*
- 7 D. Windows 2000 needs at least 64MB of RAM, although 128MB of RAM is the recommended minimum. Look over *"Hardware requirements."*
- 8 C. The `/udf:<id>,<file>` switch allows you to provide unique information on a per-computer basis. Study *"Performing Unattended Installations of Windows."*
- 9 C. The first time you log on after installing Windows, it prompts you for product activation. Each copy of Windows is validated through activation with Microsoft. Refer to *"Performing a Windows XP attended installation."*
- 10 D. The recommended RAM requirements for Windows XP are 256MB of RAM. Examine *"Hardware requirements."*

Chapter 4: Managing Files and Directories

Exam Objectives

- ✓ Working with files and directory structures
- ✓ Understanding file naming conventions
- ✓ Comprehending file attributes

Most people agree that money management is important: When you put your money somewhere, it's good to have an organized system so that you can remember where the money is and be able to retrieve it as needed. Similarly, file management is an important part of any operating system. You need a way to organize the data on your drive so that it is easily retrievable, as well as a way to identify files that are used for certain purposes.

With the changes that the Microsoft-supported file systems have undergone over the years, the A+ Certified Professional needs to have an understanding of all the file systems because it is highly likely that the computers the professional is called to work on will not all be running the most current operating system. In many cases, the computers that require recovery will be running the oldest operating systems or will have users on the bleeding edge, and trying to make use of the newest features of the latest version of Windows.

In this chapter, you get a brief history of file-naming conventions and how they have changed (and, in some cases, stayed the same) over the years. I also tell you what you need to know about file identification and file management and give you an overview of the file attributes associated with the various Microsoft file systems that have come and gone since the days of MS-DOS 1.0.

Identifying File-Naming Conventions

File-naming conventions have undergone several changes over the years. In this section, you get a look at where they have been and where they are now. You also take a look at the differences between filenames and directory entries.

When MS-DOS was the premier operating system (OS) on the market, it set the standard of what is referred to as the *eight dot three (8.3)-character filename*. This file system was created by using 32 character fields (bytes) for the filenames. If you work out the math, you're probably wondering where the extra characters are. Well, 8 plus 3 is 11, and that leaves 21 outstanding characters (bytes). Table 4-1 summarizes how each byte is used.



The File Allocation Table (FAT) is an index on your file system that holds a pointer to where each file is stored on your hard drive. On an NTFS file system, this is called the *Master File Table (MFT)*. To move a file on a file system, you only need to change its reference in the table, and not its location in the file system.

Table 4-1 Directory Entry Format

<i>Use</i>	<i>Size</i>
Filename	8 bytes
Extension	3 bytes
Attribute	1 byte
Reserved	10 bytes — FAT 32 uses two of these bytes.
Time	2 bytes
Date	2 bytes
First Cluster	2 bytes
Size	4 bytes



With every system or component used in a computer, there is a limit; sometimes it is a very large limit, but still limited. That limit is always based on a binary number. Computer systems have RAM limits of 2, 4, 8, or 16 GB, never 7.5 GB, since it is not a binary number. The binary number system is also referred to as Base 2, as opposed to our traditional number system which is Base 10. Common Base 2 numbers are determined by their digits, for instance, a 4 digit Base two number would have a maximum value of 1111 or $8+4+2+1$ or 15, and if they include a 0000 as possible value, then there are 16 possible values from 0 to 15. When you count the number of possible values for 1 digit to 12 digit Base 2 numbers, the value limits you end up with are 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048, and 4096. In cases where the limit is not based on a binary number, then there is usually some overhead involved. NetBIOS computer names are 15 characters long; but they fit into the binary system, as there is a 16th character that is used to identify the type of name, and 16 is a common binary limit. Filenames are 8 plus 3 characters long, which is only 11 characters or bytes in length, but the name is only part of a directory entry, and the entire entry is 32 bytes long.

When dealing with file attributes, one byte (8 bits) stores all attribute settings. (For more information on file attributes, see the section “Understanding File Attributes,” later in this chapter.) There are not, however, eight attributes.

The 8.3-character naming convention quickly proved to be very limiting. Extensions were often used to allow for longer names and were not directly associated with individual applications. When Windows came onto the market, people began using file extensions properly. Within Windows, the extension was used to identify the application that created a particular document.

Other OSes placed different limits on the length of filenames. Macintosh (before OS/2) limited the length of filenames to 31 characters, and OS/2 limited them to 254 characters. While 8.3 was very limiting, 254 characters seems excessive, but you likely won’t be limited in your choice of filenames.

Long and short filenames

Windows 9x and Windows XP–based OSes both allow for long filenames on a FAT16 partition. These filenames are limited to 255 characters, but some applications like Windows 2000 Explorer (199 characters, including the period separator for the file extension), Windows XP Explorer (220 characters), and Windows 98 Explorer (235 characters) cannot display all of the characters of the filename. If you want to use names longer than Windows Explorer can display, you will have to use a different application.

On a FAT16 or FAT32 partition, you are still limited to the 8.3 naming convention. Windows 9x and Windows XP-based OSes get around this problem by cheating the file system. When you save a file, it is saved using one directory entry and a short 8.3-character filename. The short filename is created by using the first 6 characters of the filename followed by a tilde (~) and an incremental number.

If you are using a Windows XP–based OS, then after creating four files with the same 6 starting characters, the formula for creating short names is changed. The first two characters are used, followed by a randomly generated 4-digit hexadecimal number, followed by a tilde (~) and the number 1.

Table 4-2 lists the names of six files and their short filenames that were created in the same directory on a Windows XP system. To get a listing of the short filenames, you can use `dir /x` or open the Properties window for each file, by right-clicking on the file and choosing Properties.

Table 4-2 Short Filenames in Windows XP

<i>Long Filename Entry</i>	<i>Short Filename Entry</i>
ShortFileTest1.txt	SHORTF~1.TXT
ShortFileTest2.txt	SHORTF~2.TXT
ShortFileTest3.txt	SHORTF~3.TXT
ShortFileTest4.txt	SHORTF~4.TXT
ShortFileTest5.txt	SH0AF2~1.TXT
ShortFileTest6.txt	SHD0C6~1.TXT

If you are using Windows 9x, then the filenames continue to increment until you run out of directory entries on a disk or you hit the limit of 65,536 entries in any given directory. As the file names increment and move to ~10 or ~100, the number of characters at the start of the name decreases to five and then four.



Microsoft has stated in different documentation for Windows 9x that it will not allow more than 99 files to be created in a directory with the same initial characters for the short filename. So if there were several files that started with the words “My File for something.txt”, then the last file that can be created in the directory has a short filename of `myfil~99.txt`. After performing tests with each version of Windows 9x, I can tell you that this stated information is wrong. This information, however wrong, has been brought up in many Microsoft exams, but is unlikely to be on CompTIA’s A+ exams.

You have now seen how short filenames are generated, but the question about where the long filenames are stored still exists. The long filenames are stored in additional empty directory entries. The characters for the long filename are stored using 11 characters per additional directory entry. So a file with a name of `My financial report for 2000.txt` takes one directory entry of the short filename (possibly `myfina~1.txt`) and one additional entry for each of the 11 characters in the filename, or an additional three entries. That means that this one file would actually occupy four directory entries on your drive. These long filename directory entries have a non-standard attribute combination of Read-only, Hidden, System, and Volume Label. Although many files on your disk may have a combination of Read-only, Hidden, and System; Volume Label is usually used alone and only on one directory entry that stores the Volume Label for the disk. By using all four of these attributes, they are a non-standard combination, and if MS-DOS systems see these entries, they ignore them rather than generating an error.

One of the problems with long filenames occurs when the long filename entries disappear. This can happen if you use MS-DOS-based disk utilities on your disk. Some of these utilities will tell you that you have a problem with your directory entries and offer to fix them. Fixing unfortunately means deleting all of the “invalid” entries, which means you lose all of your long filenames. This is not a good thing.



Microsoft makes conflicting claims about the compatibility of MS-DOS 6.x versions of `scandisk.exe` and `defrag.exe`. Some of their documentation states that MS-DOS 6.x utilities will not harm the long filename entries, while other documentation states that you should not use any MS-DOS-based file utilities. Anytime I have used MS-DOS 6.x versions of `defrag.exe` on my disks, I have lost the long filename entries, which makes me think that these utilities are not compatible with the long filename entries.

For Windows 9x, Microsoft provides a utility called either `lfnbk.exe` or `sulfnbk.exe`, depending on your version. This program runs with one of two switches, either `/b` or `/r`. The first switch backs up all of your long filename entries into a file on the root of your drive (`lfn.dat`). It also strips the current names from your file system so that older utilities can be run. After using your utilities, you can then use `lfnbk /r` to restore your long filenames to their original state.



The brunt of many a virus hoax, `sulfnbk.exe` is a valid Windows application and not a virus, as has been often misreported on the Internet.

There may come a time when you attempt to copy files to a destination that does not support long filenames. This used to happen with NetWare 3.x servers that had not enabled long filename support (OS/2 namespace). If this happens, you see a Rename File dialog box, similar to the one in Figure 4-1, for each file that has a long filename.

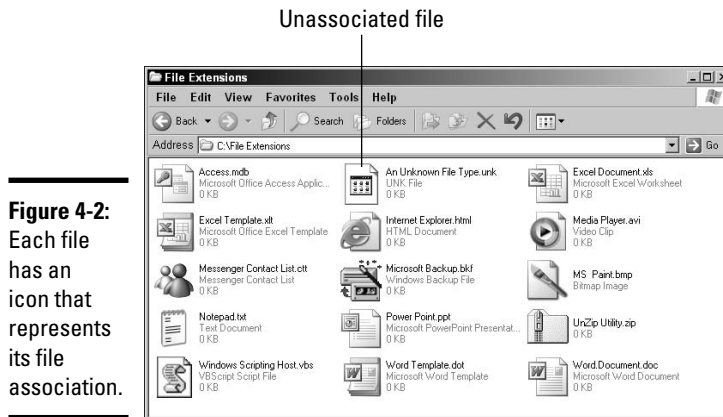
Figure 4-1: When copying files, you may be prompted to rename files if the destination does not support long filenames.



Creating file associations

Generally speaking, an *association* is a link between two separate things. *File associations* link the file extensions of filenames to the programs that can be used to open the files. When dealing with Windows, file associations let the

OS know what program should be used to open a specific type of file. This information is stored in the Registry, but can be accessed from a number of locations. The Registry is a central location that stores configuration and settings for Windows XP. Any files that are not associated with an application will have a generic Windows icon (see Figure 4-2).



Before OS X, Macintosh systems did not use file extensions for associations. Instead, each file was saved with both a data and a resource portion. The resource portion contained the file type and associated application information. Windows-based files are saved only with a data portion. In OS X, files no longer contain a resource portion, and use file extensions to define types of files.

There are a number of different ways to associate file types with applications:

- ◆ **By double-clicking unassociated files and assigning an application in the dialog boxes that follow**
- ◆ **By pressing Shift while right-clicking on associated files and selecting an Open With... option**
- ◆ **By using the Folder Options dialog box**
- ◆ **By editing the Registry**



When working with Windows 2000 or newer systems, you will not usually need to hold down the Shift key to get the Open With options. Windows XP will also display a nested menu with recently chosen options for that file type, making it easier to switch between applications for that file type.



Since there are ways other than the Registry, editing the Registry is not the recommended way. Microsoft does not recommend editing the Registry directly with the Registry editor, even though it is required from time to time. If you make a mistake with the Registry, you may have to reinstall the operating system. For more information about editing the Registry with *regedit.exe* and *regedt32.exe*, consult Book VI, Chapter 4.

For any files that are not currently associated with an application, you can double-click the file. Windows checks its list of associations; when it can not locate one, it prompts you for the name of the application that should be used, as shown in Figure 4-3. This dialog box gives you the option to Always Use the Selected Program to Open This Kind of File, which, when selected, places an association with the file extension in the Registry. This is by far one of the easiest ways to create file associations. This process is easier when using Windows XP, which also offers to check with Microsoft to see what program can be used with that file extension, in addition to letting you choose it manually, as shown in Figure 4-4.




If you are not sure about which program you should use, when you are at the Open With dialog box, you can clear the Always Use the Selected Program to Open This Kind of File check box (refer to Figure 4-3) prior to choosing the application you want to use with the file. The file will be opened by the application you chose, but the association will not be recorded in the Registry, so if you chose the wrong application, you can just double-click the file and choose another application.



Figure 4-3:
You can create file associations in the Open With dialog box.

Figure 4-4:
Windows XP offers to locate the application required to work with a file based on its extension.



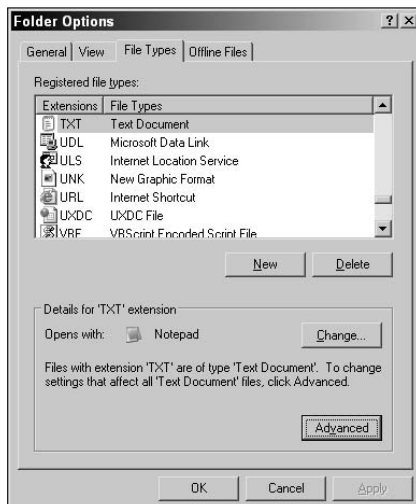
If you make a mistake or need to change the application that opens a file, you can easily do this through the shortcut menu. First select the file (if you don't do this, then the Open With option will be missing) and then right-click the icon. In the shortcut menu, you should see an option for Open With. If you are using Windows 2000, when you choose Open With, you will see the same dialog box you saw with unassociated files (refer to Figure 4-3), while if you are using Windows XP, Open With will be a nested menu, and you will also have to select Choose Program. Once again, by selecting Always Use the Selected Program to Open This Kind of File, you will be able to associate the file extension with the application you select this time through the dialog box.

If you're looking for a few more options, check out the Folder Options dialog. To get to Folder Options, follow these steps:

- 1. Open any directory window, such as My Computer.**
- 2. Open the Folder Options dialog box:**
 - If you're using Windows 9x, choose View→Folder Options.
 - If you're using Windows 2000 or newer, choose Tools→Folder Options.
- 3. When the Folder Options dialog box opens, click the File Types tab to see all of the file associations currently in the Registry.**

Windows XP makes this dialog box easy to work with by actually listing all of the file extensions in the top pane, as shown in Figure 4-5. To change the program that is associated with an application, click the Change button. This opens the Open With dialog box shown previously in Figure 4-3. From there, you choose a new application for that extension. The dialog layout and use is a little more confusing with earlier versions of Windows.

Figure 4-5:
Folder
Options in
Windows
XP.



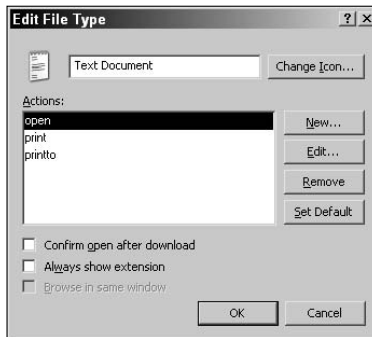
If you select a file type and click the Advanced button shown previously in Figure 4-5, you will see the Edit File Type dialog box featured in Figure 4-6. This dialog box allows you to set additional options for this type of file. These options are located under the Actions section of the dialog box and include

- ◆ **Confirm Open after Download:** This option causes a dialog box to appear after using Internet Explorer to download files with this extension. The dialog will confirm that you want to open the file.
- ◆ **Always Show Extension:** Even when Explorer is set up to hide extensions, if this option is selected, the extension will still appear for files of this type.
- ◆ **Browse in Same Window:** When loading these files into Internet Explorer, this lets Internet Explorer know if it should open a new window to show the file or if it should show the document in the current default window.

If you select one of the listed Actions, you will be able to modify how the action is performed on the selected file by clicking the Edit button and modifying the dialog which can be seen in Figure 4-7. In doing so, you will be able to specify the full path to the program to open this file type along with any switches that should be used for that particular program. For example, if the extension has open, print, and printto actions, as shown in Figure 4-7, then the field Application used to perform action for the open action may look like this:

```
C:\WINNT\system32\notepad.exe %1
```


Figure 4-6: Advanced options allow you to always have the extension of certain files displayed.



In this example, %1 is used by Windows to represent the name of the file that you are working with. When working with the print action, the line may look like this:

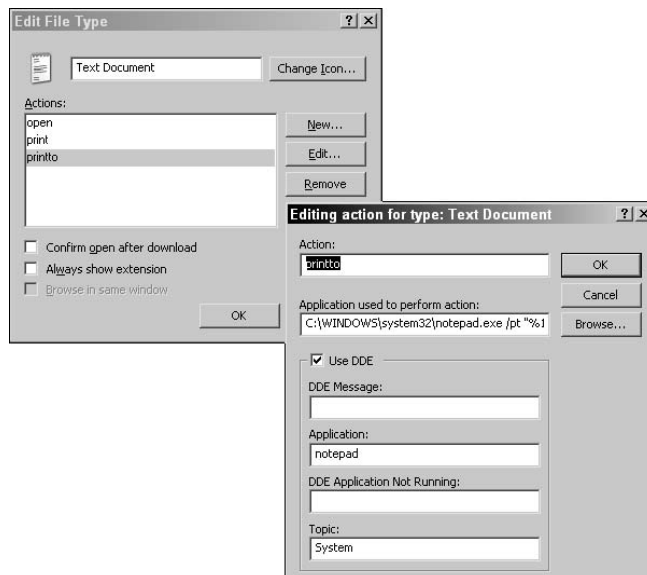
```
C:\WINNT\system32\notepad.exe /p %1
```

Finally, for the printto action, your line might look like this:

```
C:\WINNT\system32\notepad.exe /pt %1
```

Most applications will support opening files from a similar command, while fewer will support printing, or printing to a specific device.

Figure 4-7: Notepad.exe files can be opened, printed, or printed to (a device) through Windows Explorer.





Not all applications support printing for the Windows Explorer shell, but for those that do, you can either right-click the file and choose Print or drag them to a printer icon. Dragging the icon to the printer will make use of the document's printto action.

Understanding file extensions

In this section, I show you some of the major types of file extensions you can expect to see within your operating system. By default, file extensions for registered file types are hidden. Registered files types are files that have extensions which already have applications associated with them. To display file extensions within Windows XP, follow these steps:

1. Open any directory window, such as My Computer.
2. Choose Tools → Folder Options.
3. Click the View tab (shown in Figure 4-8).
4. Clear the check box in front of Hide Extensions for Known File Types.
5. Click OK.

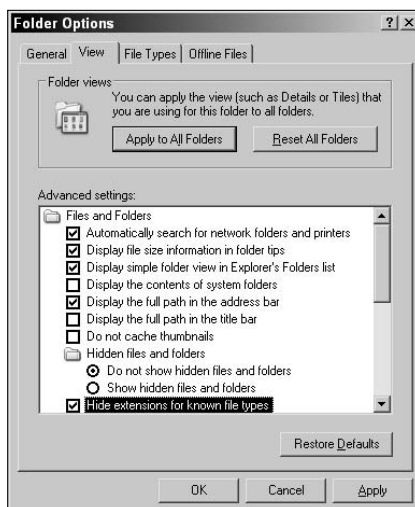


Figure 4-8:
You can tweak your Folder Options to display file extensions on your files.

In the sections that follow, I discuss common file extensions and how they are used on your computer.

Executables

Executable files can perform tasks on the system. Table 4-3 summarizes some of the file extensions common to executable files.

<i>Extension</i>	<i>Type</i>	<i>Description</i>
.bat	Batch file	Batch files are a series of commands that have been sequentially typed into a text file.
.cmd	OS/2 command file	Command files usually execute with only a command shell interface.
.com	MS-DOS command file	Like the OS/2 command files, these files usually execute with only a command shell interface.
.exe	Command-line or graphical program	This is the most common extension for executable files.
.vbs	Visual Basic Script file	Visual Basic Script files are not executable themselves, but rather require either <code>wscript.exe</code> or <code>cscript.exe</code> to be executed. They are a cross between Visual Basic applications and batch files.

Major Office applications

You need to be familiar with the file extensions associated with a few of the most popular office applications on the market. Some of these extensions are listed in Table 4-4.

<i>Extension</i>	<i>Type</i>	<i>Description</i>
.doc	Document files	Document files are usually associated with Microsoft Word, but are sometimes used by installations of WordPerfect.
.dot	Document templates	Document templates for Microsoft Word. <code>normal.dot</code> is the default template for new documents.
.ppt, .pps	Microsoft PowerPoint documents	.ppt files open in the PowerPoint interface and are immediately editable. .pps files, when you double-click them, open directly in PowerPoint Slide Show mode and close as soon as the slide show finishes.
.wks, .wk4	Lotus 1-2-3 worksheet files	The number at the end denotes that a specific version of Lotus 1-2-3 was used.

<i>Extension</i>	<i>Type</i>	<i>Description</i>
.wpd	WordPerfect document	This and .doc are common extensions for documents.
.wpg	WordPerfect Graphic	WordPerfect's proprietary graphic format.
.xls	Microsoft Excel document	Spreadsheets associated with Microsoft Excel.

Compression utilities

Over the years, and across operating systems, file compression has always been an issue, leading to the creation of numerous different compression formats. Table 4-5 lists some of the common extensions for compressed files and where they are used. Although some utilities are based in one OS or another, today you can find compression and decompression tools for most formats in most operating systems.

Table 4-5

Compression File Extensions

<i>Extension</i>	<i>Type</i>	<i>Description</i>
.ace	Ace or WinAce	A new, high-compression format.
.arc	Arc or Archive	Traditional Linux and Unix utility.
.arj	Archiver Robert Jung — the creator of the format	Traditional archive utility for MS-DOS.
.bhx, .hqx	Bin-Hex files	Bin-Hex is a Macintosh encoding format.
.cab	Cabinet file	A Microsoft format for distributing software.
.rar	Roshal Archive – created by Eugene Roshal	New, high-compression format that is used on most operating systems.
.sit	StuffIT	The most popular Macintosh compression format.
.tar	Tape ARchive files	Traditional and highly used archive format for Linux and Unix. Short for “Tape Archive,” it can back up directory structures into a single file, rather than compressing them. It is often used in conjunction with <i>gzip</i> .
.tgz	Tar-GZip	Traditional and highly used compression format for Linux and Unix. This is actually a GZIPped TAR file.
.uu, .uue	UUEncode	UUEncode is one of the most popular encoding algorithms for Linux and Unix.
.zip	PK-ZIP	This is the grandfather of PC compression formats and is still the most used format for the MS-DOS and Windows worlds.

System files

Several types of files drive your computer's operating system. Some of the extensions you should expect to see on system files are listed in Table 4-6.

Table 4-6		System File Extensions
<i>Extension</i>	<i>Type</i>	<i>Description</i>
.386	Windows-based driver files	Holds Windows formatted binary drivers.
.ani	Animated cursor files	Contains animated mouse pointers.
.bak	Backup files	Used as a default extension by a number of applications.
.bin	Binary drivers	Used for a few system drivers like <code>drvspace.bin</code> .
.cpl	Control Panel files	Contains interface to adjust Windows settings
.cur	Cursor files	Holds the graphical information to display your mouse cursor.
.da0	Backed-up .dat files	Used with the Registry files <code>system.dat</code> and <code>user.dat</code> .
.dat	Data files	Used with the Registry files <code>system.dat</code> and <code>user.dat</code> .
.dll	Dynamic Link Library	Contain common and reusable code that can be called by any application to reduce the amount of code that has to be placed within specific applications.
.drv	MS-DOS or real-mode driver files	Contains MS-DOS-formatted binary drivers
.ini	Initialization files	Text files that contain settings for applications; they are being phased out in favor of the Registry.
.msc	Microsoft Console settings files	You will see this a lot with Windows 2000 or above and a bit with Windows NT 4.0.
.msg	Error message files	Contain messages that are suppose to be displayed in the event of an error. These were used with the MS-DOS network client, and were only required if you wanted to see a text description of error codes.
.msi	Microsoft Installer files	A new file format that is used to distribute software to computers that are running the Windows Installer service.
.scr	Windows screen saver files	Holds binary code to display Windows-based screen saver
.sys	System driver files	Used with MS-DOS and Windows for backward compatibility.

<i>Extension</i>	<i>Type</i>	<i>Description</i>
.ttf	True Type Font definition files	Contains formula description used to display scalable fonts
.vxd	Virtual device drivers	Holds binary code for Windows-based device drivers

Graphic files

Graphics have been an important part of computers since computers were invented and were able to print. Early graphics were generated using the standard ASCII character set and used to create simple graphs or rudimentary graphics. A number of different formats have become popular, and with the popularity of the Internet, the number of major graphic formats has greatly increased. For a listing of graphics formats, see Table 4-7.

Table 4-7 **Graphic File Extensions**

<i>Extension</i>	<i>Type</i>	<i>Description</i>
.bmp	Bitmap	Used by MS Paint.
.eps	Adobe Encapsulated PostScript	Used by several major graphics applications.
.gif	Graphic Interchange Format	Originally owned by CompuServe, but now by AOL. It was designed to minimize download times.
.jpg, .jpeg	Joint Photographic Experts Group file format	A compression format that discards data that is thought to be invisible to the human eye. Discarding data to compress files is referred to as "lossy."
.pcd	Kodak Photo CD format	A special format of image to be displayed on devices that support Kodak Photo CD
.pcx	PC Paintbrush file	An early Windows graphic format.
.pdf	Adobe's Portable Document Format	Readable by Adobe Acrobat and Adobe Reader.
.png	Portable Network Graphics	A fairly new and open (free to use) photo standard. It is planned as a replacement for .gif and .jpeg formats; it supports all the benefits of both formats. It uses a zip type algorithm for compression.
.tif, .tiff	Aldus Tagged-Image File Format	This was used as a common interchange format between most graphics applications. It is very popular with photographic manipulators. It supports LZW (Lempel-Ziv-Welch) compression format to reduce file sizes. LZW compression is "lossless" — it doesn't discard any data.

Other file extensions

While some of the file extensions have been categorized in this chapter, Table 4-8 has some other extensions that are worth noting.

<i>Extension</i>	<i>Type</i>	<i>Description</i>
.asp, .aspx	Active Server Page files	Server-side scripted HTML files.
.cda	CD Audio files	Stores audio data on audio CDs.
.css	Cascading Style Sheet files	Stores HTML style data.
.htm, .html	HyperText Markup Language files	Store data that will be displayed on Web sites. Web servers use the HyperText Transfer Protocol (<code>http</code>) to send the information to your Web browser. <code>.html</code> is gaining popularity for displaying information locally as well. Many vendors send help files in this format so that they can be used on both their Web sites and as offline files.
.tmp	Temporary file	A short term storage file.
.txt	Text file	ASCII text files that are readable on every operating system.
.wav	Windows Audio files	Audio files designed specifically for use with Windows.



You should be familiar with all of these extensions and the categories to which they belong. Pay additional attention to executable, application, and system files, and

Lab 4-1 provides some practice working with file associations. You will require a computer with either Windows 2000 or Windows XP. Lab 4-1 can be found in the `Labs.pdf` file in the `Author` directory of the CD-ROM.

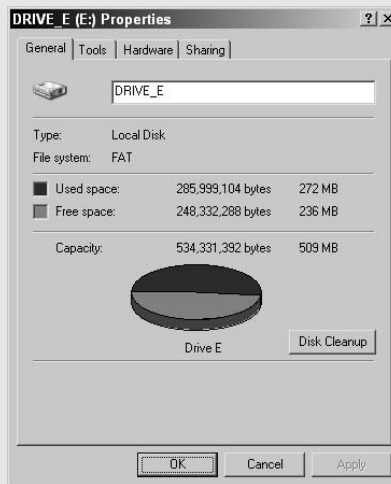
Understanding File Attributes

Since the first versions of MS-DOS and the FAT file system (see Book II, Chapter 5), files have had some basic attributes. The following sections deal with only the attributes that a user of a computer will manage, which are Read-only, Hidden, System, and Archive. The discussion focuses on the basic attributes, the extended attributes offered by Windows 2000 and Windows XP NTFS, and how to change the file attributes.

Working with the Volume Label

In addition to these four basic attributes, one more attribute, Volume, is usually set on your hard drive. The Volume attribute is usually applied to only one directory entry per drive (excluding the directory entries that are used for long filename entries) and stores the volume label for the drive. The directory entry that contains the label can be modified through one of the following methods:

- By using the `label` command at a command prompt
- By editing the Label field in the Properties window of your drive (see the figure below)
- By using the Rename command in the My Computer window (Windows 2000 and Windows XP only)



How attributes are stored

Attributes are stored as a single 8-bit binary number. Because each bit can hold either a 0 or 1, this yields 256 possible combinations of attributes — from eight 0s (no attributes applied; equal to 0 in decimal notation) to eight 1s (all attributes applied; equal to 255 in decimal notation).

Because each attribute holds a specific bit position, if several attributes are applied to a file, you can add the values (either binary or decimal) together to get a unique number that is 255 (11111111 in binary) or less for each unique combination of attributes. For example, if the file had the attributes of Read-only (1), System (4), and Archive (32), then the value of the attribute byte would be $1 + 4 + 32$, or 37. Only 37 (00100101 in binary) could represent a file with these attributes. The full list of attributes is listed in Table 4-9.

Table 4-9		File Attributes	
<i>Bit position</i>	<i>Decimal Value</i>	<i>Binary Value</i>	<i>Attribute</i>
1	1	00000001	Read-only
2	2	00000010	Hidden
3	4	00000100	System
4	8	00001000	Volume Label
5	16	00010000	Subdirectory
6	32	00100000	Archive
7	64	01000000	Unused
8	128	10000000	Unused



For the exam, you will be asked only about the Read-only, Hidden, System, and Archive attributes. But you should be aware of the Volume Label and Subdirectory attributes for real-world applications. You do not need to know the bit values of any of the attributes, just know what the four basic attributes are used for.

The basic attributes

All files on your computer, regardless of your Microsoft operating system or file system, have these four basic attributes:

- ◆ **Read-only**
- ◆ **Hidden**
- ◆ **System**
- ◆ **Archive**

Descriptions of the four common attributes are listed in Table 4-10.

Table 4-10		Basic File Attributes
<i>Attribute</i>	<i>Description</i>	
Read-only	Files with this attribute cannot be deleted from the command line and cannot be written to or saved over.	
Hidden	These files are not visible to users unless those users have changed their viewing options to display hidden files. Hidden files cannot have their attributes changed by the <code>attrib.exe</code> command.	
System	System files are flagged as being required by the operating system. They cannot have their attributes changed by the <code>attrib.exe</code> command.	

<i>Attribute</i>	<i>Description</i>
Archive	Files with the archive attribute are ready for archiving or backing up. This attribute is used by some backup programs so that they can perform incremental backups. Currently, the <code>xcopy</code> command supports the <code>/a</code> switch, indicating that only files that have the Archive attribute set should be copied. When the file is modified, the Archive attribute will automatically be reset on it.

Any time files are modified, the Archive attribute is automatically set, so after Windows XP is installed, all files will have the Archive attribute enabled. Windows XP also sets Read-only, Hidden, and System attributes on many required system files, to prevent accidental modification or deletion by the user; but other than these files, the only attribute that is automatically set is the Archive attribute. You will see how to manually set these attributes later in this chapter when you read “Setting basic attributes”.

Windows 2000 and Windows XP extended attributes

Windows 2000 and other newer Microsoft operating systems have three other file attributes that are used on NTFS partitions. These three attributes are often referred to as *extended attributes*:

- ◆ **Compress:** Used to enable on-the-fly file level compression of data
- ◆ **Encrypt:** Used to enable on-the-fly file level encryption of data
- ◆ **Index:** Used to allow files to be indexed by the OS Indexing Service.

The sections that follow take a closer look at each of these attributes.

Compress

The Compress attribute allows and restricts file compression. To compress a file or folder, follow these steps:

- 1. Right-click the file or folder and choose Properties.**

The Properties dialog box appears.

- 2. Click the Advanced button.**

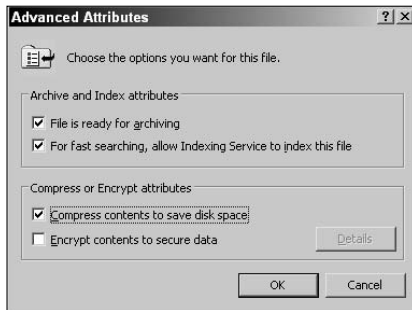
The Advanced Attributes dialog box pops up.

- 3. You will be able to compress the file or folder by selecting the Compress Contents to Save Disk Space check box, as shown in Figure 4-9.**

If you are asked about extended attributes on your exam, look at the OS that is mentioned in the question. If the question deals with Windows NT 4.0 rather than Windows 2000 or Windows XP, remember that Windows NT 4.0 supports the Compress attribute but not the Encrypt or Index attributes.



Figure 4-9:
The
Compress
attribute is
set on a file-
by-file or
folder basis.



To compress a file, then file needs to be read in to memory, where compression takes place, and then re-written to the disk is it's compressed form. When you change the compress file attribute, then the OS saves the change to attribute prior to actually going to steps of compression the file. This is done in two steps, since the attributes are stored in the MFT (Master File Table), and not in the actual file itself. After the attribute change is recorded, file compression for that file is performed. There is a small chance that something may go wrong during the compression phase, so that some of the files on your hard drive may have the Compress attribute set but may not actually be compressed. One situation that may cause this inconsistent Compress attribute is a power interruption during the compression process, so that the attribute change was made but the compressed version of the file was not saved. To ensure that this has not happened, you could use the command line utility, `compact.exe`.

`compact.exe` can be used to compress or decompress files from the command line. It supports switches that will allow it to process subdirectories as well. If you suspect inconsistent compression states for your files, you can use the `/F` or force switch, which goes through the process of compressing all files, even if they already have the compress attribute, this does not double compress files, but rather ensures that all files that are suppose the be compressed actually are compressed. See Table 4-11 for a complete list of the switches that `compact.exe` makes use of. If you run `compact.exe` without any switches, it will display compression information for the files in your current directory.

Table 4-11 `compact.exe` command switches

<i>Switch</i>	<i>Description</i>
<code>/C</code>	Compresses the specified files. Directories will be marked so that files added afterward will be compressed.
<code>/U</code>	Uncompresses the specified files. Directories will be marked so that files added afterward will not be compressed.

<i>Switch</i>	<i>Description</i>
/S	Performs the specified operation on files in the given directory and all subdirectories. Default "dir" is the current directory.
/A	Displays files with the hidden or system attributes. These files are omitted by default.
/I	Continues performing the specified operation even after errors have occurred. By default, COMPACT stops when an error is encountered.
/F	Forces the compress operation on all specified files, even those which are already compressed. Already-compressed files are skipped by default.
/Q	Reports only the most essential information.



If you move a compressed file within a partition, the file will retain its Compression attribute. However, if you copy a compressed file to a folder, it will inherit the compression attribute that is set for the destination folder.

Encrypt

An encrypted file is protected only against being read. This means that the file can be moved to another location on the same partition or renamed; these actions require modifying the directory table and do not constitute reading the file. The file may also be deleted if the appropriate NTFS permissions have not been applied to it.

To encrypt a file or folder, follow these steps:

1. Right-click the file or folder and choose Properties.

The Properties dialog box appears.

2. Click the Advanced button.

The Advanced Attributes dialog box pops up.

3. You will be able to encrypt the file or folder by selecting the Encrypt Contents to Secure Data check box



If the user who encrypted the file, moves or copies a file between NTFS partitions (even between computers), then the file will remain encrypted. This is different from Compression or NTFS permissions. If the file is copied to a non-NTFS partition, then the file is automatically decrypted. If you are not able to decrypt the file, then you will receive an *Access Denied* message when you attempt to move the files to a location which is on a different partition.



PKI (Public Key Infrastructure) is a system that allows for authentication users and encryption of data. PKI works with key pairs that are used for in conjunction to encrypt or decrypt data. The pair of keys is defined as private (known only to the user) and public (known to all other users of the infrastructure). Public keys are usually stored in certificates. When data is

encrypted by using a public key, it can only be decrypted by using the private key. EFS makes use of PKI; if you do not have an infrastructure, the local workstation will auto-generate keys on the local computer to support EFS. If you are using EFS within an organization, you should make use of an enterprise-wide PKI, which can be implemented by using Microsoft's Windows Server 2003 Certificate Authority.

Windows does not allow you to both encrypt and compress a file. This is because compression requires that the file be rewritten, and encryption does not allow the file to be rewritten. The reason encryption does not allow the file to be rewritten is because the file header contains the file's encryption keys. These keys are protected by the public key of the user who encrypted the file as well as the public key of the Encrypted File System (EFS) Recovery Agents. The EFS Recovery Agent is specified in the local public key policies of the computer or in the active directory. By default on a workstation or Active Directory domain, the EFS Recovery Agent will be set to the Administrator account. When a file is encrypted, the only people who can read the file are the user who enabled encryption, the EFS Recovery Agent, and any other users who have specifically had their public keys used to encrypt the file's encryption keys.

To allow other people to access your encrypted files, follow these steps for a file that has been encrypted:

1. Right-click the file or folder and choose Properties.

The Properties dialog box appears.

2. Click the Advanced button.

The Advanced Attributes dialog box pops up.

3. Click the Details button.

The Encryption Details for the file dialog box pops up. This will list all users who can transparently access the file.

4. Click the Add button.

The Select User dialog box pops up.

5. If the user you want to grant access to is not on the list, click the Find User button to access the standard OS Select User dialog box; otherwise, select the user you want to add and click the OK button.

6. Click the OK button on each of the other dialogs that you had opened.

If you do not have an enterprise PKI, you will not be able to access certificates for other users; so you will not be able to easily grant other users access to encrypted files.



Similar to `compress.exe` for compressing files, you can use the command-line utility `cipher.exe` to encrypt files from the command line or in batch files.

In the same way that you could use `compress.exe` for compressed files, you can use `cipher.exe` for encrypted files. Table 4-12 reviews the switches that can be used with the `cipher.exe` command.

Table 4-12 **cipher.exe command switches**

<i>Switch</i>	<i>Description</i>
/A	Operates on files as well as directories. The encrypted file could become decrypted when it is modified if the parent directory is not encrypted. It is recommended that you encrypt the file and the parent directory.
/D	Decrypts the specified directories. Directories will be marked so that files added afterward will not be encrypted.
/E	Encrypts the specified directories. Directories will be marked so that files added afterward will be encrypted.
/F	Forces the encryption operation on all specified objects, even those which are already encrypted. Already-encrypted objects are skipped by default.
/H	Displays files with the hidden or system attributes. These files are omitted by default.
/I	Continues performing the specified operation even after errors have occurred. By default, CIPHER stops when an error is encountered.
/K	Creates new file encryption key for the user running CIPHER. If this option is chosen, all the other options will be ignored.
/N	This option only works with /U. This will prevent keys being updated. This is used to find all the encrypted files on the local drives.
/Q	Reports only the most essential information.
/R	Generates an EFS recovery agent key and certificate, then writes them to a .PFX file (containing certificate and private key) and a .CER file (containing only the certificate). An administrator may add the contents of the .CER to the EFS recovery policy to create the recovery agent for users, and import the .PFX to recover individual files.
/S	Performs the specified operation on directories in the given directory and all subdirectories.
/U	Tries to touch all the encrypted files on local drives. This will update user's file encryption key or recovery agent's key to the current ones if they are changed. This option does not work with other options except /N.
/W	Removes data from available unused disk space on the entire volume. If this option is chosen, all other options are ignored. The directory specified can be anywhere in a local volume. If it is a mount point or points to a directory in another volume, the data on that volume will be removed.
/X	Backup EFS certificate and keys into file filename. If <code>efsfile</code> is provided, the current user's certificate(s) used to encrypt the file will be backed up. Otherwise, the user's current EFS certificate and keys will be backed up.

Index

Indexing the file system will greatly improve the time it takes to search your hard drive, but in return, indexing demands space on your drive. Indexing will not only record the file names and locations, but will also record words that are found in your files. The indexing service has a set of noise words that it ignores, such as *it*, *that*, *is*, *not*, and *the*. As files are indexed, the results are stored in a catalog, and when searching for files you are able to specify the catalog you want to use. By default there is a catalog configured to index your local files on your hard drives. This default catalog is called *system* and is found in the System Volume Information directory of your boot drive or C: drive.

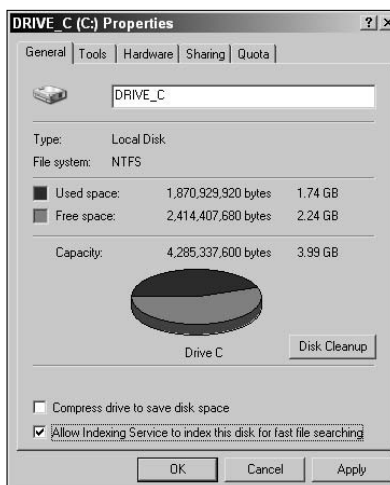
To enable basic indexing of your drive, follow these steps:

1. Right-click your drive in My Computer and select Properties (see Figure 4-10).
2. Select the check box labeled For Fast Searching, Allow Indexing Service to Index This file.

This enables indexing, provided that the Indexing Service is running.

You can verify that the Indexing Service is running by using the Services Administrative Tool, Start→Control Panel→Administrative Tools→Services, and verifying that the Indexing Service is started.

Figure 4-10: Indexing can be enabled for your whole drive or for just a folder.



To ensure that the Indexing Service is running on Windows XP, do the following:

1. Choose Start⇨Search.
2. In the Search companion pane on the left, click Change Preferences.
3. In the results that appear, click With Indexing Service (For Faster Local Searches).

If the results show an option for Without Indexing Service, then you are already set to use them. By following either of these links, the left pane of the search window allows you to change the Indexing Service setting (see Figure 4-11).

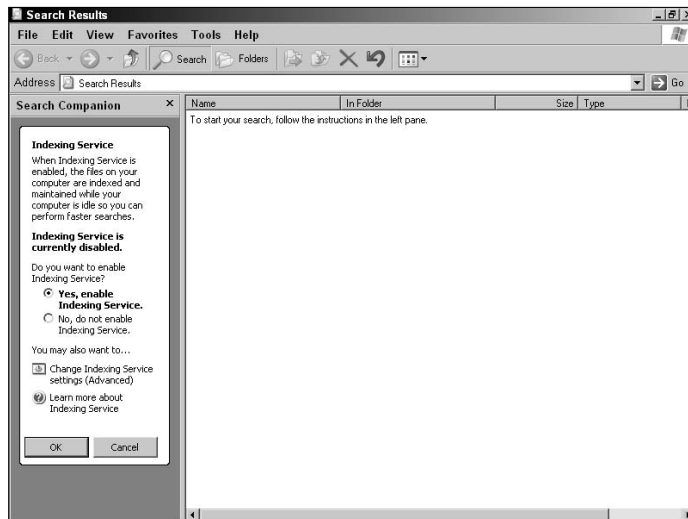


Figure 4-11:
In order for indexing to work, it must be enabled.

Setting basic attributes

As with so many other areas of the operating system, where there are many different ways to achieve your goal, there are many different ways to change the attributes of files. The first method — using the command line — is the only way that is common to all of the Microsoft operating systems; but we will also show you how to set attributes using the Windows GUI.

Command line

In order to change file attributes from the command line, you will have to get to the Command Prompt window by choosing Start⇨All Programs⇨

Accessories⇒Command Prompt, which will open a Command Prompt window. To set attributes you have to use the `attrib.exe` command. The basic syntax for the command is this:

```
attrib <attribute to set (H, S, R, A)> <files to modify>
```

Setting of the attributes includes a plus (+) or minus (-) sign, specifying whether the attribute should be enabled (a plus) or disabled (a minus), followed by a one-letter code for each attribute. For example, to set the read only attribute and remove the archive attribute from the `C:\boot.ini` file, you would type, `attrib.exe +R -A C:\boot.ini`. The one-letter codes are as follows:

- ◆ H: Hidden
- ◆ S: System
- ◆ R: Read-only
- ◆ A: Archive

If you use the `attrib` command without any options, it will display a list of the current attributes of all the files in the current directory. Here is a sample of what that would look like:

```
C:\>attrib.exe
A      C:\7flen.crt
A SHR  C:\arclldr.exe
A SHR  C:\arcsetup.exe
A      C:\AUTOEXEC.BAT
A SH   C:\boot.ini
A      C:\CONFIG.SYS
A SHR  C:\IO.SYS
A SHR  C:\MSDOS.SYS
A SHR  C:\NTBOOTDD.SYS
A SHR  C:\NTDETECT.COM
A SHR  C:\ntldr
A      C:\uninst.log
```

If you use the `attrib.exe` command with options but without specifying a filename, you will set the attribute(s) for all the files in the current directory. Here is an example of the command and the attributes that would be set:

```
C:\>attrib.exe -A -S -H -R

      C:\7flen.crt
      C:\arclldr.exe
      C:\arcsetup.exe
      C:\AUTOEXEC.BAT
      C:\boot.ini
      C:\CONFIG.SYS
```

```
C:\IO.SYS
C:\MSDOS.SYS
C:\NTBOOTDD.SYS
C:\NTDETECT.COM
C:\ntldr
C:\uninst.log
```

You can specify a file after the `attrib` command to see the current attributes of that file, or you can include the one-letter codes to set the attribute(s), like this:

```
C:\>attrib.exe C:\config.sys
A      C:\CONFIG.SYS
C:\>attrib.exe +h +r C:\config.sys
C:\>attrib.exe C:\config.sys
A HR   C:\CONFIG.SYS
```

Another option you can include with the `attrib.exe` command is `/s`, which instructs `attrib.exe` to process the files in all subdirectories — this does not change the attribute on the subdirectories, just the files in the subdirectories. If you are using Windows 2000 or newer Windows OS, then `/d` will change the attributes of the directories as well.

There are some restrictions on changing attributes with the `attrib.exe` command. Hidden and System files will not have their attributes changed, unless the current Hidden or System attribute is also specified in a single command. Examine the following command sequence, which shows an example of dealing with Hidden or System files.

```
C:\>attrib.exe -S -H -A -R C:\config.sys

C:\>attrib.exe C:\config.sys
      C:\config.sys

C:\>attrib.exe +S C:\config.sys

C:\>attrib.exe +R C:\config.sys
Not resetting system file - C:\config.sys

C:\>attrib.exe C:\config.sys
      S      C:\config.sys

C:\>attrib.exe +R +S C:\config.sys

C:\>attrib.exe C:\config.sys
      S R    C:\config.sys

C:\>attrib.exe +H C:\config.sys
Not resetting system file - C:\config.sys
```

```
C:\>attrib.exe +H +S C:\config.sys

C:\>attrib.exe C:\config.sys
SHR C:\config.sys

C:\>attrib.exe -S -H -R C:\config.sys

C:\>attrib.exe C:\config.sys
C:\config.sys
```

Windows GUI

In addition to the command line interface for changing attributes, Windows offers a GUI interface. When you right-click a file, you only see the options for Read-only, Hidden, and Archive when working with FAT partitions. If you are working with NTFS partitions, then the Archive check box is replaced with an Advanced button, as shown in Figure 4-12.

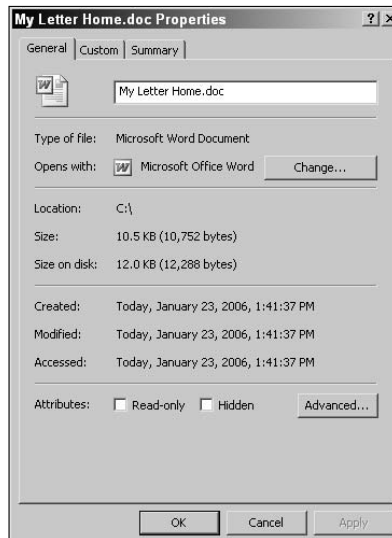


Figure 4-12:
On NTFS
partitions,
you see only
the Read-
only and
Hidden
attributes.

Clicking the Advanced button lets you set the Archive attribute as well as the Windows 2000/XP advanced attributes of Index, Compress, and Encrypt. See Figure 4-13.



Lab 4-2 provides some practice working with attributes via the command line. For this exercise, you need a computer running Windows 2000 or Windows XP. Lab 4-2 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Figure 4-13:
The
Advanced
Attribute
window
is only
available
on NTFS
partitions.



Getting an A+

This chapter discusses file and directory management from the point of view of file naming and file attributes. The following major points are covered:

- ◆ Windows XP is not limited to the 8.3 naming convention.
- ◆ Windows XP supports simulated 8.3 naming for backward compatibility.
- ◆ The four basic file attributes are Read-only, Hidden, System, and Archive.
- ◆ Windows 2000 and newer Windows OSs support additional file attributes for Compression, Encryption, and Indexing. These attributes are only available on the Windows NT File System (NTFS).
- ◆ The Archive attribute is still used by many backup programs to locate files that have changed since the last backup.

Prep Test

- 1 What is the maximum character length of a filename under Windows XP?
 - A 64
 - B 128
 - C 254
 - D 255
 - E 256

- 2 You get a call from Billy. You find out that he has a problem with several of the files on his hard drive. He claims that Windows XP has renamed many of his files and directories. After questioning him, you find out that he just defragmented his hard drive with an MS-DOS 6.0 boot disk and the MS-DOS version of `defrag.exe`. What is likely the cause of the problem, and is there a solution to correct it?
 - A The MS-DOS version of `defrag.exe` was unaware of the long filename directory entries and overwrote the directory entries. There is no easy fix for this problem.
 - B The long filename entries have just become unassociated from the files. Run `assoc c: /repair` to re-establish the associations.
 - C The filename entries are fine; they were automatically backed up, and Billy just needs to run `1fnbk` from the root of each drive.
 - D Everything is fine; Billy just needs to reboot.

- 3 Tom is copying files to a NetWare server on his network, and he is prompted to rename his file. What is likely the reason?
 - A The NetWare server already has a file with the same name.
 - B Tom does not have sufficient permissions on the server.
 - C The server does not support long filenames.
 - D The server needs to be rebooted.

- 4 Which of the following can be used to associate file extensions with an application? Choose all that apply.
 - A `attrib.exe`
 - B `assoc.exe`
 - C Shift+right-click and choose Open With
 - D Tools → Folder Options → Associate tab

- 5** Which of the following are extensions of executable files? Choose all that apply.
- A .ini
 - B .prg
 - C .exe
 - D .cmd
- 6** Which of the following file extensions are associated with compression programs? Choose all that apply.
- A .com
 - B .zip
 - C .cmp
 - D .cab
- 7** Which of the following attributes are associated with files on your hard drive? Choose all that apply.
- A Restrict-Modify
 - B Archive
 - C Visible
 - D Hidden
- 8** When you copy an NTFS compressed file to a FAT32 partition, what happens to the file?
- A The copy action is denied.
 - B The file is copied, but remains compressed, so it is unreadable.
 - C The file is copied but becomes uncompressed.
 - D The file can be copied only if the FAT32 partition has been compressed with DriveSpace compression.
- 9** You are moving several of Bob's files from a directory to another directory on the same partition. What message should you receive when you encounter an encrypted file?
- A You should not receive a message; the file will just be moved.
 - B You should get an access denied message.
 - C You should get a message stating that you could not read the file, and that it may be corrupted.
 - D You should get a message stating that the file cannot be moved because it is encrypted.

- 10** You are moving several of Bob's files from a directory on an NTFS partition to a directory on a FAT32 partition. What message should you receive when you encounter an encrypted file?
- A You should not receive a message; the file will just be moved.
 - B You should get an access denied message.
 - C You should get a message stating that you could not read the file, and that it may be corrupted.
 - D You should get a message stating that the file cannot be moved because it is encrypted.

Answers

- 1** D. The maximum length of a filename is 255 characters. See *“Long and short filenames.”*
- 2** A. You should only use Windows 2000 or Windows XP utilities to maintain your disks. Review *“Long and short filenames.”*
- 3** C. When copying files to destinations that do not support long filenames, you are prompted to rename the files. Check out *“Long and short filenames.”*
- 4** C. `attrib.exe` is used to change file attributes and not file associations. `assoc.exe` does not exist, and the correct path on a Windows XP computer is `Tools→Folder Options→File Types`. Peruse *“Creating file associations.”*
- 5** C, D. `.ini` is an initialization file, and `.prg` is a made-up extension. `.cmd` files are text files similar to batch files but are treated as executable. Take a look at *“Understanding file extensions.”*
- 6** B, D. Pk-zip files use the `.zip` extension, and `.cab` files are a proprietary Microsoft compression format. Peek at *“Understanding file extensions.”*
- 7** B, D. The four basic attributes are Read-only, Hidden, System, and Archive. The advanced NTFS attributes are Compress, Encrypt, and Index. Look over *“Understanding File Attributes.”*
- 8** C. Compressed files copied from NTFS volumes to any other type of disk will become uncompressed. If they are copied to other NTFS volumes, they will inherit the compression attribute from the destination folder. Study *“Windows 2000 and Windows XP extended attributes.”*
- 9** A. As long as the move is within one partition, then the files will just be moved. If the directories are on different drives, then you would receive an access denied message. Refer to *“Windows 2000 and Windows XP extended attributes.”*
- 10** B. Since the files are being moved to a different drive, and you are not the person who encrypted the files, you will be denied access. If you had owned the files, the files would have been decrypted as they were copied. Examine *“Windows 2000 and Windows XP extended attributes.”*

Chapter 5: Command Prompt Procedures

Exam Objectives

- ✓ **Identifying the fundamentals of using command-line utilities such as `cmd.exe`, `format.com`, and `attrib.exe` to manage Windows 2000 and Windows XP**
- ✓ **Making use of `edit.com` to create and edit a text file**
- ✓ **Using command-line tools such as `md`, `cd`, `rd`, `dir`, `copy`, and `xcopy.exe` to manage directories and files**

In this chapter, you take a look at several command prompt utilities and a few graphical utilities. Even in the age of Windows, statements that can be issued from command prompt are still required knowledge for support professionals. This knowledge will help you automate processes — such as batch files — and solve problems when the graphical operating system is not functioning. After reading this chapter, you should have a good understanding of many of the basic commands that are available from the command prompt.

As an added benefit, you find out how to build a basic batch file with some controls in it. I put this information at the end of the chapter because you won't be tested directly on that knowledge, but it will sure come in handy on the job.

As an A+ Certified Professional, you will want to use the full range of tools that are available to you. This chapter introduces you to many of the command-line tools that allow you to quickly diagnose and repair problems with operating systems.

Using `command.com` and `cmd.exe`

Just as many utilities in Windows 9x are command-line-based, many of the really useful utilities in Windows 2000 and Windows XP are command-line-based. Windows 2000 and Windows XP give you two options for running command-line utilities: `command.com` and `cmd.exe`.

When you worked with MS-DOS and Windows 9x, the file `command.com` was the command-line interpreter, and it served as the basic method of executing programs in the OS. Some of the programs that were executed were programs internal to `command.com`, while others were external and could be found on the hard drive.

Windows NT-based OS, like Windows 2000 and Windows XP, use `cmd.exe` as the basic command-line interpreter for issuing commands. When using `cmd.exe`, you are using a Windows 32-bit application from the point of view of memory management and application stability. `command.com` still exists on these operating systems, but it exists for backward compatibility for older applications.

Typically, when you want to use a command prompt within Windows XP, you should run `cmd.exe` by choosing Start→Run, typing `cmd.exe` and clicking OK. `cmd.exe` can also be found in the Start menu under All Programs→Accessories→Command Prompt. This runs as a 32-bit Windows application and thereby can spawn other applications or support standard memory management. Each copy of `cmd.exe` that you execute appears on the process list in Task Manager. If you run `cmd.exe`, the first two lines tell you what has executed:

```
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.
```

`command.com` on the other hand, is a 16-bit MS-DOS application. It loads a simulated 16-bit environment in the form of `ntvdm.exe` (NT Virtual DOS Machine). You will see `ntvdm.exe` listed on the process tab in Task Manager. The first two lines that will be listed in this window are:

```
Microsoft (R) Windows DOS
(C) Copyright Microsoft Corp 1990-2001.
```

You will notice that, in this case, you are running DOS and Windows XP. You will also notice that your long filenames are no longer supported. If you use the command prompt to launch other 16-bit Windows applications, then they will be launched within the NTVDM that is executing. The *NTVDM* is the environment that all 16-bit MS-DOS applications and 16-bit Windows applications run in. For more information about supporting these applications, refer to Book VI, Chapter 2.

`cmd.exe` is used when running all other applications. It is a 32-bit command prompt with a number of enhancements, including:

- ◆ The ability to run 16-bit and 32-bit command-line utilities
- ◆ The ability to run Posix and OS applications
- ◆ Built-in doskey support for saving and executing commands



Windows 2000 and Windows XP use `cmd.exe` as the main command-line interface and provide `command.com` for backward compatibility.

Managing Directories

All files that are saved on a disk are stored in a hierarchical directory structure. All the files could be placed at the top of this structure, but it would be disorganized, and therefore poorly managed. Also, most *top-level directories*, or *root directories*, can hold a limited number of files. This directory structure can hold as many nested subdirectories as you want. The commands discussed in the following sections let you get listings of files as well as create and delete directories on your disk.

dir

The first command that you will see is the directory command (`dir`). This command is used to get a listing of the files that are in a directory on your disk. Typing `dir` by itself gives you the listing of your current directory. The current directory is usually listed in the command prompt, like this:

```
C:\WINDOWS\COMMAND>_
```

The `dir` command is very useful and has several options that are discussed later in this section. First, though, you need a firm grasp of wildcards, relative paths, and absolute paths.

Two wildcards can be used to modify what results you are given back: `*` and `?`. The `*` represents one or more characters. Here is an example of its usage to retrieve a list of files that match a certain pattern. The command

```
dir WIN*XE
```

returns the following:

```
WINMINE.EXE
WINHELP.EXE
WINHLP32.EXE
WININIT.EXE
WINVER.EXE
WINFILE.EXE
WINPOPOP.EXE
WINIPCFG.EXE
WINREP.EXE
```

The results include all files that start with `win` and end with `xe`, with any number of characters in between — even zero. These wildcards are also useful with copy commands, which I discuss later in the chapter.

The `?` works a little different than the `*` because it represents one or zero characters. Here is an example of the `?` in action, as it looks for all files that were created on the fifth of any month after Jan 01, 2000. The command

```
dir ??052???.TXT
```

would return the following list of files:

```
1052000.TXT
5052000.TXT
10052000.TXT
11052000.TXT
12052000.TXT
4052002.TXT
12052003.TXT
6052008.TXT
```

The `?` in the command doesn't return any files that have more than two characters before the 05 and doesn't return any more than three characters after the 2 in the `dir` statement. This is often helpful when files in a directory are named with six- or eight-character numeric dates with the pattern of `mmd-dyyyy.txt`, and you are looking for all of the files that are named for the fifth of any month.

When you type `dir`, you are given a directory listing for your current directory. If you want to see the listing for a different directory, Table 5-1 provides several options for choosing a different directory. All the command examples in the table use the `c:\parent_dir\child_dir\grandchild_dir` directory structure, with the current directory set to `c:\parent_dir\child_dir`.

Table 5-1 Ways to Specify Directory Paths

<i>Command</i>	<i>Directory Result</i>
<code>dir</code>	Returns the listing for the current directory, <code>c:\parent_dir\child_dir</code> .
<code>dir c:</code>	Returns the listing for the current directory on the C drive <code>c:\parent_dir\child_dir</code> .
<code>dir c:\parent_dir</code>	Returns the listing for <code>c:\parent_dir</code> .
<code>dir \</code>	Returns the listing for the root of the current drive <code>c:\</code> .
<code>dir ..</code>	Returns the listing for the parent directory of the current directory <code>c:\parent_dir</code> .
<code>dir ..\..</code>	Returns the listing of the parent directory of the parent directory of the current directory <code>c:\..</code> .
<code>dir ..\child_dir</code>	Returns the listing of a directory named <code>child_dir</code> , which is a child of the parent directory <code>c:\parent_dir\child_dir</code> .

Command	Directory Result
<code>dir grandchild_dir</code>	Returns the listing of a child directory named <code>grandchild_dir</code> , <code>c:\parent_dir\child_dir\grandchild_dir</code> .

Whenever a full path is specified, starting with the drive letter, it is referred to as an *absolute path*. If you do not specify the full path, you are using a *relative path*. Care should always be taken when using relative paths with commands. Look, for example, at the following code sample:

```
dir sub_dir
del *.*
```

In this example, a directory listing is taken of a subdirectory, while the following delete command was working on the entire contents of the current directory. This is an easy mistake to make. When working with relative paths, the double period (`..`) notation refers to a parent directory, and the single period (`.`) refers to the current directory.

Most programs and commands can have their actions modified by providing options on the command line. These options are usually represented by one or more letters and are introduced to the command with either `/` or `-`. The `-` is usually used with commands that come from (converted from) the UNIX OS, while most MS-DOS programs have implemented the `/`. These options are referred to as *switches*. `dir /on /s` is an example of the `dir` command using switches.

Table 5-2 summarizes some of the most important switches for the `dir` command.

Switch	Description
<code>/ax</code>	The <code>/a</code> switch is short for “attributes.” This switch provides a listing of files that have matching attributes. This switch must be used in conjunction with an additional letter to provide results. There are five letters that may be used: (d)irectory, (a)rchive, (h)idden, (s)ystem, and (r)ead-only. Using a minus sign with any five of the letters to reverse the listing, for example, <code>/a-d</code> , shows you things that are not directories.
<code>/b</code>	The <code>/b</code> command displays a bare listing. The bare listing does not include a separate section in the output that tells you the directory that you are working with, but rather displays a single-line listing, like this: <pre>c:\windows\hosts.txt c:\windows\lmhosts.txt</pre>

(continued)

Table 5-2 (continued)

<i>Switch</i>	<i>Description</i>
<code>/ox</code>	The <code>/o</code> switch is short for “order by.” This switch is similar to <code>/a</code> in that <code>/o</code> by itself will not work; it requires an additional letter to tell it how to order or sort. The options that are available for sorting are: (n)ame, (s)ize, (e)xtension, (a)ccessed date — earliest first — and (d)ate modified — earliest first. If you use a minus sign (for example, <code>/o-d</code>), the order is reversed. If you use the letter <code>g</code> after the <code>o</code> (for example, <code>/ogd</code>), then directories will be grouped at the top of the list rather than mixed in.
<code>/p</code>	This switch pauses the screen after each full screen of text and waits for a key to be pressed. If you do not use this option, you can use the pipe-more command, which looks like this: <code>dir c:\windows*.exe more.</code>
<code>/s</code>	This switch will include listings for each subdirectory under in the directory listing.
<code>/w</code>	This switch will display the text in a wide listing. It enables more text to be displayed on a screen by using multiple columns.
<code>/x</code>	This switch is used on Windows 2000 and Windows XP computers to display the short filenames, as well as the long filenames.

If you want to use the same set of switches each time you execute the `dir` command, you can use the `dircmd` environment variable. This variable can be set in the `autoexec.bat` file or at a command prompt, like this:

```
set dircmd=/on /w
```

Now, every time you type `dir`, you will get output that looks like this:

```
Volume in drive D is DRIVE_D
Volume Serial Number is FFFF-FFFF

Directory of D:\DOS

[.]          [..]          CHOICE.COM   DELTREE.EXE  DOSHELP.HLP
DOSKEY.COM   EXPAND.EXE    FC.EXE       FDISK.EXE    HELP.COM
HELP.HLP     HIMEM.SYS     INTERLNK.EXE INTERSVR.EXE LOADFIX.COM
SMARTDRV.EXE TREE.COM
              15 File(s)      525,943 bytes
              2 Dir(s)    446,431,232 bytes free
```

You should note that the list is presented in wide format and sorted by name.

If you are using Windows 2000 or Windows XP and you want the setting of the `dircmd` variable to be saved, then you should add the variable by using the Environment Variables settings in the System Control Panel, which you can get to using `Start → Control Panels → System`, to open the System Properties dialog box; then click the Advanced tab and the Environment Variables button.

mkdir

`mkdir`, or `md`, is used to create directories; there is no difference between the two commands other than their spelling. Many of the commands that are used with MS-DOS originated in other operating systems, and in some cases, new short forms were created to make them easier to use. In an effort to provide backward compatibility, support for the older spelling of the command was also kept. The directory created will be in the current directory unless you provide an alternative path to the command, like this:

```
mkdir "c:\temp\my new directory"
```

chdir

`chdir`, or `cd`, is used to change the current directory for a drive to another directory. The drive need not be your current drive; this command can set a current directory on another drive. For example, if your current drive is C:, you could still type `cd d:\my_dir_on_d` to change the current directory for the D: drive. You will not see a difference on your screen, but if you change to the D drive by typing `d:`, you will see that the current directory is set to `d:\my_dir_on_d`. The current directory is important when you want to use other file operation commands, such as `copy`. If you only specify the drive that you want to work with, then you will be working with the current directory on that drive. Take a look at this example:

```
C:\Documents and Settings\ed>c:  
C:\>cd \  
C:\>mkdir d:\old_configs  
C:\>cd d:\old_configs  
C:\>copy a*.bat d:  
C:\>copy c:\c*.sys d:\old_configs
```

The first line changes your current drive to C:. The second changes to the root directory of the current drive (c:\). The third creates a new directory on the D: drive, and the fourth line sets `old_configs` as the current directory on the D: drive. You will see the `copy` command in the section “copy,” later in this chapter, but the fifth line copies all files in the current directory that start with `a` and end with `.bat` to the current directory on the D: drive (currently `old_configs`). The last line copies all files from `c:\` that start with `c` and end with `.sys` to `d:\old_configs`. Lines 5 and 6 copy files from the same directory (c:\), while the relative path is used in line 5 and the absolute path is used in line 6. The destination directory in line 6 is also absolute, while the relative path or current directory is used in line 5. In both statements (lines 5 and 6), the same directories are used as the source and destination directories. Relative paths can save on typing, but they can also cause errors if you are not careful.

rmdir

`rmdir`, or `rd`, is used to remove or delete directories from your drive. Two rules are imposed on you: Before you delete a directory, the directory must be empty, and it cannot be the current directory. You can empty a directory by using the `del` command to delete the files. To remove a directory, just specify its location after the `rmdir` command:

```
rmdir c:\remove_me
```

If you are using Windows 2000 or Windows XP, there is an optional switch, `/s`, that will automatically delete subdirectories and files.



All of the command-line tools in this section have not changed substantially since MS-DOS.

Copying and Moving Files

Doesn't it seem like when you finally get things organized, it's time to start all over again? When organizing and backing up files, you will often be required to copy or move files to new locations, either in another directory or on another disk. The following sections provide an overview of the commands that let you do this.

copy

The `copy` command expects you to give at least the name of the file you would like to copy. If you provide only one filename, then the selected file is copied into the current directory. If you provide a source filename and a destination directory by using a command like

```
copy c:\source\myfile.txt c:\destination
```

then the file will be copied into the destination directory. You can also rename files while you are copying them by using a command like this:

```
copy c:\source\*.bat c:\destination\*.old
```

The previous command would copy all of the files with an extension of `.bat` from the source directory and rename them with an `.old` extension in the destination directory.

If you are about to overwrite an existing file, you will be prompted to confirm the operation. This can be suppressed if you use `/Y` at the end of your copy command. `/Y` answers "yes" to the copy command's confirm overwrite

questions. If you want to consistently overwrite destination files, you can set the `copycmd` environment variable to `/Y` in the same manner that you set the `dircmd` variable in the “dir” section, earlier in this chapter.

xcopy

Many times, you will have to copy entire directory structures from one location to another. If you were to do this with the `copy` command, you would first have to create all of the destination directories by using the `md` command. With the `xcopy` command, you can perform this task in a minimal amount of time. To copy an existing directory named `source` to a new directory named `destination`, you would type the following command:

```
xcopy c:\source\*. * c:\destination\*. *
```

If you wanted to copy all the subdirectories as well, you would use:

```
xcopy c:\source\*. * c:\destination\*. * /s
```

To also include empty directories, add `/e` (empty) to the end of the command. To include just files with the archive attribute set, add the `/a` (archive) switch.

Like the `copy` command, adding `/Y` will overwrite files without asking for confirmation. The `/Y` tells the command to answer “yes” to all overwrite prompts.

move

The `move` command moves files from one directory to another. It is also used to rename directories. To use the `move` command, you have to specify the name of the files you want to move and then specify the destination directory, like this:

```
move c:\source\source_file.txt c:\destination\
```

This example moves the file `source_file.txt` into the directory `c:\destination\`. If the destination directory does not exist, then you will see an error message.

If you want to rename the directory `c:\source`, you use:

```
move c:\source destination
```

del

To delete or remove files or directories, you can use `del` or `erase`. Once again, these two commands are synonymous. If you want to delete multiple

files, you have to use the * and ? wildcards. MS-DOS and Windows 9x let you delete only files with this command, while Windows 2000 and Windows XP will also delete directories.



It is very easy to unintentionally delete files (or delete the wrong files) when working with relative paths. When possible, use full pathnames to avoid mistakes.

ren

The `ren` command is used to rename files and directories. Similar to many of the commands that you have looked at, you specify the source name and a new name for the file or directory.



All of the command-line tools in this section have not changed substantially since MS-DOS.

Making Comparisons

There are a many special function commands, and this section will look at a couple commands that can be used to compare files. These commands are `fc.exe` and `diskcomp.exe`.

diskcopy.com and diskcomp.com

These two tools do not get very much use these days. Both of these tools were designed to be used with floppy disks and both take two parameters, first is the drive letter of the source disk and the second is the destination disk. Both of these can be the same drive letter, such as `diskcopy a: a:`, which will copy the contents of drive `a:` to a temporary location and then ask for the target disk to be inserted in the `a:` drive. `diskcopy.com` will copy the contents of the disk to a new disk, and `diskcomp.com` will compare the contents of the two disks to ensure that there are no errors during or after the copy.

fc.exe

`fc.exe` is still a valid and useful tool, it compares the contents of two files and indicates the differences between the two files. There are a variety of switches with this command that will allow you to modify how the output looks or the way that the comparisons are made. The basic format of the command only requires that you specify the two files that you want to compare using a format like `fc.exe c:\boot.ini c:\old-boot.ini`. Here is a sample of the `fc.exe` command in action.

```
C:\>fc.exe boot.ini old-boot.ini
Comparing files boot.ini and OLD-BOOT.INI
***** boot.ini
[boot loader]
timeout=30
default=multi(0)disk(0)rdisk(0)partition(1)\WINDOWS
***** OLD-BOOT.INI
[boot loader]
timeout=20
default=multi(0)disk(0)rdisk(0)partition(1)\WINDOWS
*****
```

In this example, there was one difference between the two files, in the `timeout=` line. By default, `fc.exe` shows the line before and after the difference in the files, and shows you which lines came from which file.

This command is very useful when attempting to compare two very long configuration files for differences.



All of the command-line tools in this section have not changed substantially since MS-DOS.



The exam objectives call for testing on command-line tools and they may not be limited to just the common tools that are listed in this book. To review new command-line tools, choose Start⇨Help and Support and search for *New Command-line Tools*. To review all possible command line tools, Choose Start⇨Help and Support, and search for *Command-line reference A-Z*.

Performing Diagnostics and Tuning Performance

At the command prompt, a few commands can be used to optimize and diagnose your computer. `defrag.exe` can be used to optimize your hard drive, while `chkdsk.exe` can be used to check your disk for errors. You would have used `mem.exe` when creating an MS-DOS boot disk, to view memory used when trying to optimize your memory usage.

The MS-DOS or command prompt versions of the `defrag.exe` and `chkdsk.exe` utilities have now been replaced by Windows XP and Windows 2000 versions.

defrag.exe

The purpose of the `defrag` utility is to fix speed and performance problems with hard drives. As files are written to and then deleted from a disk, they leave holes or blank areas scattered around your drive. When you write files

to a disk, they always write to the largest open spaces that are available to them. There will be times when the largest area is not large enough for the entire file, and the file will have to be split into pieces. These *fragmented* files are slower to access because the disk head has to keep moving to a new location on the drive. If you want to defragment the files on your hard drive, you can use Computer Management in Windows 2000 or Windows XP, or `defrag` in Windows XP. In an effort to provide command-line access to all Windows-based utilities, Microsoft provides command-line access to the graphical defragmenter utility.

In either case, you are using the defragmentation utility from Executive Software, and you have the option to analyze only or to force defragmentation when disk space is low.

For more information about using Computer Management in Windows 2000 and Windows XP, see Book VI, Chapter 4.



`Defrag.exe` does not exist for Windows 2000. It is only available on Windows XP and newer Windows operating systems.

chkdsk.exe

`chkdsk.exe` serves an important purpose within the Windows operating system. Its job is to check the directory structure and directory entries for corruption, and the disk for worn-out areas. Its purpose is to reduce the chance of data loss by catching corruption early and by fixing small problems before they become larger.



For more information about `chkdsk.exe`, refer to Book VI, Chapter 4.

mem.exe

The `mem` command displays usage information about your computer's memory. This is often useful when trying to optimize the boot process or when trying to maximize the amount of memory available to MS-DOS-based applications. `mem` runs from within MS-DOS.

Running the command by itself displays basic information about how much memory is available to the MS-DOS environment. Two main switches are available to the `mem` command:

- ◆ `/c` (classify) tells you how much space applications are using in upper memory and conventional memory.
- ◆ `/d` (debug) gives you a detailed breakdown of what is stored in the first 1MB of RAM.

Working with the Rest

The rest of the utilities in this chapter defy easy classification and are presented in the following sections. Although they serve a wide range of functions, they are all important in your computer's usage.

attrib.exe

All files have four basic attributes: read-only (r), hidden (h), system (s), and archive (a). The `attrib` command lets you change these attributes. The attributes are added or removed from files by specifying the attribute with a + or - in front of the filename, as in the following statement:

```
attrib -s -h +a +r c:\*.sys
```

The previous statement removes the system and hidden attributes while adding the archive and read-only attributes. If `/s` were added to the end of the line, then it would have been done to all of the files in the subdirectories as well.

As the name suggests, *read-only* files cannot normally be deleted, nor can they be modified. *Hidden* files are not usually visible. *System* files have special file protection so that you may not delete or modify them. *Archive* files are files that have been modified. The archive attribute is used by some backup utilities to identify changed files for incremental backups.

diskpart.exe

To manage disk partitioning, if the need ever arises, Windows 2000 and Windows XP use a graphical disk partitioning tool called Disk Management; Windows XP also uses another command-line tool called `diskpart.exe`. Although this tool is capable of all disk partitioning tasks, you use this tool only if you need to perform rare disk partition changes that Disk Management cannot perform, such as expanding a partition on a basic disk. The reason for not using `diskpart.exe` is that its user interface is awkward. To illustrate the interface, the following is the set of steps required to create a partition on a disk and format the drive (extra blank lines have been removed):

```
C:\>diskpart.exe
DISKPART>list disk
  Disk ###  Status         Size      Free      Dyn  Gpt
  -----  -
  Disk 0    Online          6142 MB     0 B
  Disk 1    Online           510 MB     0 B
  Disk 2    Online          6142 MB   6142 MB

DISKPART> select disk 2
```

```
Disk 2 is now the selected disk.

DISKPART> create partition primary size=1000
DiskPart succeeded in creating the specified partition.

DISKPART> list partition
  Partition ###  Type              Size      Offset
  -----
  Partition 1    Primary           1004 MB   32 KB

DISKPART> select partition 1
Partition 1 is now the selected partition.

DISKPART> assign letter=f
DiskPart successfully assigned the drive letter or mount point.

DISKPART> exit
Leaving DiskPart...

C:\>format f: /fs:ntfs
The type of the file system is RAW.
The new file system is NTFS.
WARNING, ALL DATA ON NON-REMOVABLE DISK
DRIVE F: WILL BE LOST!
Proceed with Format (Y/N)? y
Verifying 1004M
Volume label (ENTER for none)? More Data
Creating file system structures.
Format complete.
    1028128 KB total disk space.
    1020600 KB are available.
```

This process could have been completed through a single wizard in Disk Management, with less effort.

format.com

The `format.com` command prepares a disk to be used by your computer. The purpose of this command is to check to see if all of the clusters on the partition are in working order and to create the directory table. On FAT partitions, the directory table is referred to as the *File Allocation Table*. The directory table maintains a list of where each file starts on the disk.

The `format` command requires a drive letter and supports additional switches. The proper syntax to format your A: drive is

```
format a:
```

You could modify the command with `/q` to perform a quick format of the drive. The quick format doesn't check the integrity of the sectors on the drive but only deletes and re-creates the directory table.

If you are using MS-DOS or Windows 9x, you can also use the `/s` (system files) switch. This copies over the system files that are required to boot.

If you are using Windows 2000 or Windows XP, you will be able to use */FS: filesystem* to specify the format of the partition as either FAT, FAT32, or NTFS.

sys.com

The MS-DOS *sys*, or system command, enables you to put the files that are required for booting your computer on a disk that has already been formatted. By default, the system files are copied from your current drive to the drive specified in the command. You can copy the files from an alternate location, but you cannot copy them to your current drive. Here are two examples:

```
sys a:
sys d:\sysfiles a:
```

Because the files have to be placed in a specific location on the disk, there is a chance that you will be told that there is no room on the disk if that area is already occupied.

ver

The *ver* command indicates what version of the command prompt you are running. This varies between copies of Windows 9x and Windows 2000 or XP, as shown in Table 5-3. *ver* with Windows 9x operating systems lets you see that they are all sub-versions of Windows 4.0, while all copies of Windows 2000 and XP are sub-versions of Windows NT 5.0. Major and minor version numbers from software manufacturers let you know how much they think the software has changed, so architecturally each of these operating systems has not changed significantly since their initial releases.

<i>Version</i>	<i>Output of ver Command</i>
Windows 95 Retail	Windows 95. [Version 4.00.0095]
Windows 95 OSR2	Windows 95. [Version 4.00.1111]
Windows 98	Windows 98. [Version 4.10.1998]
Windows 98 SE	Windows 98. [Version 4.10.2154]
Windows 2000 Professional	Microsoft Windows 2000 [Version 5.00.2195]
Windows XP Professional	Microsoft Windows XP [Version 5.1.2600]
Windows Server 2003	Microsoft Windows [Version 5.2.3790]

help

If you need to know how to do something with the command line, you have two main options. All commands support */?* as a switch to get additional

information on how to use the command. In addition to this feature, Windows XP also has a `help` command. To use `help`, type `help` followed by the command that you would like to get help on — for example:

```
help help
```

If you type `help` by itself, you are given a list of commands that you can request help for.

Working with the edit Command

Most of the configuration files on your computer can be edited with any text editor. If you have the Windows interface loaded, you will probably prefer to use a graphical program like Notepad or Wordpad. In the event that you don't have a computer that is currently running Windows, you can always use the `edit` command.

Opening files

To open files, you first have to launch the editor. This can be done by typing `edit` in any command window. And if you know the name of the file that you want to edit, you can specify it after the `edit` command, like this:

```
edit a:\config.sys
```

If you have the editor open and you want to open a file, you have to press `Alt` to get access to the menu, press `F` to access the File menu, and then press `O` to open. This will bring you to the Open dialog box, shown in Figure 5-1. By default, you will be in the File Name position, with the current name of `*.*` selected. If you type in the full path and filename of a file, you will be able to open the file. If you choose to browse for your file, it gets a little more complicated.

Figure 5-1: At first glance, navigating the MS-DOS-based dialog boxes can be confusing.



There are two ways to navigate the Open dialog window. First, you can use the Alt key on the keyboard. You will notice white letters in several places in the dialog box. By pressing Alt and a white letter, you will be able to go to a section of the dialog box. For example, pressing Alt+D moves the cursor to the Directories pane, and pressing Alt+F moves it to the Files pane. You can use the arrow keys on your keyboard to scroll through either of these lists and use the Enter key to select choices. The first entry in the Directories pane is . . to navigate to the parent directory, and the last entries in that pane are all of your drive letters, to switch to different drives.

The second way to navigate from pane to pane in this dialog box is to use the Tab key. If you press the Tab key in succession, you will move through the panes, fields, and buttons from left to right, top to bottom. The exact order for this dialog would be: File Name, Files, Directories, Open Read-only, Open Binary, OK, Cancel, Help. By holding the Shift key while tabbing, you move in reverse order. To select the check boxes, press the spacebar when the option is highlighted.

When the correct file has been selected, you can press Enter or select the OK button to open the file.

Saving files

Saving files also requires that you first get access to the menu bar by pressing the Alt key. Alt+F opens the file menu, Alt+S saves, and Alt+A opens the Save As dialog box. If the file has already been saved, then saving updates the original; but if the file has never been saved, then either saving option opens the Save As dialog box, shown in Figure 5-2. This dialog box is similar to the Open dialog box with regard to navigation. After you select a directory and provide a name, select the OK button.

Figure 5-2:
When saving, choose a location and then a filename.



Searching and replacing

One of the more powerful features of `edit` is its ability to search and replace text in your file. To perform a search and replace, press the `Alt` key to access the menu, and then press `S` to access the Search menu. The Replace dialog box (shown in Figure 5-3) can be opened by pressing `R`. You can then type the string that you are looking for, as well as the string that you would like to replace it with.

Figure 5-3: Search and replace is a powerful feature.



To narrow down the search, you can also tell the editor to search only whole words and to match the case. Searching for whole words means that when you search for *cat*, for example, you will not find it as part of the words *cats* or *categories*. Matching case means that *cat* would not be found if you were searching for *Cat*.

Closing the editor

To close the editor, press the `Alt` key; then press `F` for File and `X` for Exit. If you have unsaved changes, you are prompted to save your changes, as shown in Figure 5-4.

Figure 5-4: If there are unsaved changes, you are prompted to save them.



Batch File Basics

You now know several basic commands that make the command prompt easier to use — as well as a powerful ally. In the following sections, you find out how to write a batch file, which makes the command prompt an even more powerful tool. A *batch file* is simply a string of commands, which could have been typed at a command prompt, chained together to automate a process.

Starting your file

To start things off, you need a place to put the commands that you want to execute, and, to do this, you need to make a new text file and save it with a `.bat` extension. Open the text editor of your choice and create a new file; then save your file with a `.bat` extension. If you're using a Windows-based editor, such as Notepad and the operating system option of Hide File Extensions for Known File Types is still enabled, you will have to enclose your filename in quotation marks when saving it.

If you want to show or hide your file extensions, use the following steps:

1. **Open My Computer.**
2. **Choose Tools⇨Folder Options.**
3. **Click the View tab.**
4. **Clear (to show) or check (to hide) the box next to Hide File Extensions for Known File Types.**
5. **Click OK.**

With your batch file open in the text editor, you can now type any commands you want, but remember to put only one command per line in your file. You can use this file to create, copy, or delete files, or anything else that you could do from the command prompt.

Once you have created a batch file, you can execute or run it by typing the path to file and the filename at the Start⇨Run command or at a command prompt.



It is best to do some testing with the batch file on a test computer — I have seen people create batch files with improperly formatted `rmdir` commands, which then deleted many more files and directories than expected. If you use commands that format partitions, remove files, or delete disk partitions, note that poorly structured commands may result in you being required to reinstall the operating system on your computer. The commands suggested in this section are not dangerous.

Getting your directions

Many people want to be able save the output of their commands, or more rarely, answer questions or provide input for a command. There are a few ways to do this, and Table 5-4 summarizes the differences among the options. All of these options require that your program is generating output going to `stdout` (standard output) or accepting input from `stdin` (standard input). `stdout` and `stdin` are interfaces or objects that C programmers use almost exclusively when requiring data input and output for command-line programs. As a programmer, I can write an application that sends its output to an interface called `stdout`, and then I let the operating system deal with where that actually is, as it might be a teletype terminal, a printer, a monitor, or a text file. For Windows XP, data that is sent to `stdout` will be displayed in a command prompt window.

Table 5-4		Redirection Symbols
<i>Symbol</i>	<i>Function</i>	
<	Sends the contents of a file listed on the right into the command on the left. This is useful when you want to provide an answer to a question that a command will prompt for (for example, when the <code>format</code> command prompts for confirmation). By creating a one-character text file with just <code>Y</code> and a return, you will be able to pass a <code>yes</code> into the <code>format</code> command.	
>	By specifying a command on the left and a text file or <code>nul</code> on the right, you can prevent output of a command from going to the screen. Instead, the output will be passed to the file that is listed on the right side. If the file exists, it will be overwritten. This is extremely useful if you want to log the results of certain statements.	
>>	Similar to the <code>></code> symbol except that the output is appended to the destination file if it already exists. If the file does not exist, it will be created.	
	Referred to as a pipe. It pipes or directs the output of the file that is listed on the right side into (as input) the command that is listed on the left. This may allow you to bypass the creation of a text file.	

Working with parameters or starting arguments

You can get a lot more mileage from your batch files if you pass and use arguments with them. An argument is a piece of data that is passed to the batch file or command when typing the command, such as `copy C:\this.txt C:\that.txt`, which has two arguments: `C:\this.txt` and `C:\that.txt`. There are ten arguments that you can use with a batch file. This really drops off to nine because the first argument is `%0`, and that represents the name of the running batch file. `%1` through to `%9` can be used as arguments for your batch file. If you have created a batch file that copies files from one directory to another, you can make your batch file more reusable by working with arguments. Here is an example of this concept:

Old batch file — `StaticCopy.bat`:

```
xcopy c:\dos\*.* c:\dosback
```

This file could be made more generic with these modifications:

New batch file — `DynamicCopy.bat`:

```
xcopy %1 %2
```

This new batch file would be run by typing:

```
dynamiccopy c:\dos\*.* c:\dosback
```

It's true that you are typing more to achieve the same result. But you have to consider that this line could be in a much larger batch file, which could use the two parameters (`%1` and `%2`) in many different ways. This could actually save you a lot of typing when creating and editing your batch file because the parameters can be used rather than hard coding the paths or parameters.

This system works fine until you find yourself in a position in which you need more than nine parameters passed to your batch file. However, if you plan the order in which you use your parameters, you can actually support many more than nine. This is done with the `shift` command. Each time you use `shift` in your batch file, all of the parameters that are passed to your batch file are shifted over one position. For example, `%2` becomes `%1`, `%3` becomes `%2`, and so on.

Making batch files make decisions

So far, you have seen some ways to make your batch file more generic. Now, I'll show you how to make it think for itself, or at least think for you. There are two ways to get your batch file to make some decisions, although after you've read this section, you may argue that there is only one way, since both ways make use of batch file labels.

First, the `if` and `:label` statements introduce simple program flow logic. For example, if you know that you need two parameters passed to your batch file for it to work properly, then you can use something like this in your batch file:

```
if not !%2==! goto correct

:Instructions
echo.
echo Please provided two parameters. The first is the
echo source directory, and the second is the destination.
echo.
```

```
goto end

:correct
xcopy %1 %2
goto end

:end
```

When working with the `if` statement, you evaluate whether two options are the same or whether a condition is true. In this case, you are testing for two items being equal, and, in order to do this, you need two items — you cannot test equality with only one item. In this example, you want to see if the second parameter is blank, so you use `%2`, which represents the second parameter; but you aren't able to use `if not %2==`, as this would evaluate `%2` to nothing, and you need something on the other side of the equation. This gets worse if the second parameter was not passed to the batch file, as `%2` will not exist, and you are then asking if nothing equates to nothing. The `if` statement cannot handle evaluating nothing in an equation. In order to make sure there is something to evaluate on both sides of the equation, and not change the results, you need to add a character to both sides of the equation. In this example, I chose to use the exclamation mark. You should also take note of the use of the word `not` to reverse the results of the test.

Some other items that the `if` statement can test for includes OS environment variables and whether files and directories exist. To test an environment variable for a specific value, you use a command like this one that checks the name of the computer that is running the batch file:

```
if !%COMPUTERNAME%==EDSPC goto instructions
```

To test if a file exists, you use something like:

```
if exist c:\autoexec.bat goto Instructions
```

If you want to know if a directory exists, you can test for `nul`, which only exists for your test if a directory is there.

```
if exist c:\windows\nul goto Instructions
```

The full example also shows you how to work with labels and the `goto` statement. The `goto` statement uses only the first eight characters of the label, so try to keep label names unique. The `goto` statement requires a destination in the batch file to be sent to, and that is where the labels come in. Although `end` is a label that is meaningful to you, to the batch file it is just another label. Remember, if you use a `goto` in your batch, you must have a matching label or you will get errors and your batch file will not work.

The last method of controlling the flow of your batch file is using the `choice` or `set` statement. The `choice` statement allows you to pause your MS-DOS or Windows 9x batch file to wait for input. Input to the `choice` statement is a single character. Here is an example of how to use the `choice` statement:

```
@echo off
choice /c:ynm You choose
if errorlevel == 3 goto m
if errorlevel == 2 goto n
if errorlevel == 1 goto y
goto end

:y
echo You chose Yes.
goto end

:n
echo You chose No.
goto end

:m
echo You chose Maybe.
goto end

:end
```

The `choice` statement uses the list of parameters after the `/c:` to listen to input. One of those choices must be provided. The `choice` statement will actually ignore any other input and wait for one of those choices.

Also, take note of how to handle the input afterward. Each choice is assigned a number, based on its order: Y = 1, N = 2, M = 3. The `if` statement should actually be read as, “If the error level is greater than or equal to x, go to label.” Because of how the computer handles `if` statements, always test the `errorlevel` from the highest to the lowest. If you test the values in the wrong order, and start you evaluation with `if errorlevel = 1 goto y`, then everything will satisfy the `this` test because all of the possible values are greater than or equal to 1.

If you are using Windows 2000 or Windows XP, then you can do the same type of thing in your batch file using the `set` command, but I have found that `set` is an inferior tool for this type of test. So I suggest getting `choice.exe` from either a Windows 98 computer, the Windows 2000 resource kit, or extracting it from Microsoft provided tools such as the Microsoft Platform Support Reporting Utility. The Platform Support Reporting Utility is available from the download page at <http://support.microsoft.com/default.aspx?scid=kb;en-us;816819> and provides `choice.exe` in most packages, including the Base/Setup/Storage/Print/Performance package.

Starting with Windows Server 2003, Microsoft has included `choice.exe`, and it supports slightly different options than the MS-DOS version. This version of `choice.exe` is also compatible with Windows XP. Here is a copy of the same batch file using the Windows Server 2003 `choice` command. Notice the `/M` and the quotes around the message text:

```
@echo off
choice /c:ynm /M "You choose"
if errorlevel == 3 goto m
if errorlevel == 2 goto n
if errorlevel == 1 goto y
goto end

:y
echo You chose Yes.
goto end

:n
echo You chose No.
goto end

:m
echo You chose Maybe.
goto end

:end
```

The `if`, `choice`, and `set` statements add many options to the design of your batch files; they can help you make batch files that are flexible and comprehensive.

Looping

I hate having to repeat myself, and if you feel the same way, then loops are for you. Loops can be used inside or outside of batch files with only a small modification. The loop is a `for` loop and requires a list of values that are assigned to a variable and a statement that uses that variable. The variable is assigned with a `%` sign in front of it outside of a batch file, and two `%` signs when used within a batch file. Here is an example of a loop at the command prompt:

```
echo off
for %x in (hello world) do echo %x
```

The output of the previous loop would be:

```
hello
world
```

If this were done in a batch file, it would look like this:

```
echo off
for %%x in (hello world) do echo %%x
```

Loops may cut down your work when you have several repetitive lines in your batch file.

Locating a command

When you type a command at a command prompt, in a batch file, or at the Run command, your computer follows a set of steps to locate the command. The steps include:

1. The computer checks the current directory for the command.
2. If the command is not found in the current directory, the computer checks all directories listed in the Path environment variable.

If the command is not located, possibly because you misspelled it, then you see the following message:

```
Bad command or file name
```

If you see this message, check the spelling of the command and try again.

The path variable can be set in the `autoexec.bat` file for all Microsoft operating systems. It is the only environment variable that can be set without using the `set` command. To set the path variable, you can use either

```
set path = c:\windows;c:\windows\system
```

or

```
path = c:\windows;c:\windows\system
```

The only advantage to the latter is that it uses four fewer keystrokes. If both `c:\windows` and `c:\windows\system` are in the path, then they will be checked when a command that is typed is not in the current directory.

Windows 2000 computers provide an additional location for setting the path environment variable. To change the path on a Windows 2000 computer, open the System Control Panel, click the Advanced tab, and click the Environment Variables button. You can use the Environment Variables dialog box to modify the path for the computer.

Putting it together

The following is a short batch file that has been documented with `rem` or remark statements throughout the file. It shows a completed file that uses many of the techniques discussed in this section. This batch file was designed to gather the TCP/IP configuration from a computer and place it in a text file on the `C:` drive; but it could place the information on a drive on a server so that it could be collected by an administrator. The text file that the batch file creates is named for the computer information from which the information was gathered. This data is collected only once. The networking troubleshooting commands used in this file can be found in Book IX, Chapter 3.

```
rem Check for existing file, and if it exists, exit the batch file,
rem and make the directory if it does not exist.
if not exist C:\IPSettings\nul mkdir C:\IPSettings
if exist C:\IPSettings\%computername%.txt goto end
rem Check to see if the computer is running Windows 9x
rem or Windows NT style OS.
if not !%OS%!=!Windows_NT goto Win9x
:WinNT
rem Create a new file containing the output of ipconfig.exe
rem which is the current TCP/IP Configuration.
ipconfig /all > C:\IPSettings\%computername%.txt
rem Add to this file the current routing table using netstat.exe
netstat.exe -R >> C:\IPSettings\%computername%.txt
rem Skip the Windows 9x section and goto the end of the file.
goto end
:Win9x
rem Use Windows 9x compatible versions of the Windows NT commands
rem networking commands.
winipcfg /all /batch C:\IPSettings\%computername%.txt
netstat.exe -R >> C:\IPSettings\%computername%.txt
:end
```

If this file is saved to `C:\Windows\GetIPSettings.bat`, it can be executed by specifying the name of the file `GetIPSettings` in a command window or by typing it after using `Start`→`Run`. By specifying the name, the entire path on your computer will be searched for the file with either `.com`, `.exe`, `.cmd`, or `.bat` as the file extension.

Getting an A+

The most important things to remember from this chapter are what the different commands are used for and the major differences between similar commands. Here is a quick review:

- ◆ `attrib` changes file attributes, such as read-only, hidden, system, and archive.
- ◆ `copy` copies one or multiple files from one directory to another.

- ◆ `defrag` defragments your hard drive to increase performance.
- ◆ `dir` displays a list of files in a directory or directories.
- ◆ `edit` modifies text or ASCII files.
- ◆ `format` prepares a disk for accepting data or to erase a disk.
- ◆ `mem` displays information about memory usage.
- ◆ `chkdsk` checks a disk for errors that could cause data loss.
- ◆ `ver` displays the version of Windows that is in use.
- ◆ `xcopy` is like `copy` but can be used to copy entire directory structures.

Prep Test

- 1** You are working on the help desk, and you receive a call from Mary. She is looking for a file that is saved on her hard drive. She has forgotten the name she gave it, but she knows that the title contains the word “budget” and that it is a Microsoft Excel spreadsheet. What command will best help her locate her document?

 - A `dir *budget.xls`
 - B `dir *budget*.xls /b`
 - C `dir *.budget*.xls /s`
 - D `dir *budget*.xls /b /s`
- 2** Sitting at your computer one day, you attempt to open a file with the `edit` command, but you see in the title bar that the file is opened as read-only. How can you open the file for editing?

 - A Open the file using `edit <filename> /modify`.
 - B Before opening the file, type `attrib -r <filename>`.
 - C After the file is open, type `attrib -r <filename>`.
 - D When you attempt to save the file, check off the Overwrite Existing File option.
- 3** You need to regularly back up files in a directory, but you want to back up only files that have changed. What commands can you use to accomplish this in the easiest way possible?

 - A `fc` and `xcopy`
 - B `attrib` and `copy`
 - C `fc`, `changes`, and `xcopy`
 - D `xcopy` and `attrib`
- 4** You want to copy a directory structure (complete with subdirectories), so you type `copy c:\dir1*. * c:\newdir /s /e`. What will you find in the `newdir` directory?

 - A All of the files that were in `dir1`
 - B All of the files and subdirectories that were in `dir1`
 - C Nothing — the command is improperly constructed and will generate an error
 - D None of the above

- 5** What is the purpose of the `ver` command?
- A To switch output from other commands to verbose mode
 - B To imitate older versions of the command shell to improve backward compatibility
 - C To display the version of the operating system that you are using
 - D To enable advanced logging options
- 6** You receive a call from a user on your network complaining that his computer is slow. You ask a few more questions and find out that he feels that the speed of his hard drive is much slower than when he got his computer. What command would you suggest running?
- A `diskfix`
 - B `quikdisk`
 - C `scanfix`
 - D `defrag`
- 7** You have created a new partition on your hard disk, and assigned it a drive letter of `F:`. The new partition needs to be formatted as NTFS. Which of the following commands will perform the required action?
- A `format.exe F: /FS:NTFS`
 - B `filesystem F: /FS:NTFS`
 - C `format.com F: /FS:NTFS`
 - D This can only be performed when creating partition using Disk Management.
- 8** You have just been told that an MS-DOS application that you are supposed to run requires 550K of conventional memory. How can you check to see whether you have that much memory available?
- A `convchck`
 - B `memchck`
 - C `mem`
 - D `freemem`
- 9** You need to copy an entire directory structure from floppy disk over to your hard drive. Which command should you use?
- A `xcopy.exe`
 - B `pathcopy.exe`
 - C `copy.exe`
 - D `dircopy.exe`

Command Prompt Procedures

10 You are using a Windows 2000 computer and want to remove the `c:\files` directory. This directory is not empty. What command do you use?

- A `deltree`
- B `rmdir`
- C `deldir`
- D `nukedir`

Answers

- 1 D. Although C looks like a good choice as well, using the subdirectory switch, it has an additional period in the search string so will not return as many files and will work only if Mary's document has a period in front of the word *budget*. See "dir."
- 2 B. The `attrib.exe` command only lets you change the attribute before the file is opened. Review "attrib.exe."
- 3 D. The `attrib.exe` command can be used to remove the archive attribute from all of the files in a directory structure. Then as files are modified, you can use `xcopy *.* c:\backuplocation /s /a` to copy all of the files to a new location. If you want your next copy to copy the files that were modified since the last copy, you will have to use the `attrib` command to remove the archive attribute again. Check out "xcopy."
- 4 C. The suggested command will cause a syntax error and nothing will be copied. The `copy` command does not support either `/s` or `/e`; in order to use these options, you have to use the `xcopy.exe` command. Peruse "xcopy."
- 5 C. `ver` displays the version of the operating system. Take a look at "ver."
- 6 D. `defrag` is used to reorganize files, which speeds up access to the disk. Peek at "defrag."
- 7 C. To assign a specific file system to a drive during formatting, you use the `/FS:` option. Look over "format.com."
- 8 C. `mem` tells you how much conventional memory is available to your applications. Refer to "mem.exe."
- 9 A. `xcopy.exe` can copy entire directory structures. Examine "xcopy."
- 10 B. `rmdir` removes directories on a Windows 2000 computer even if they contain files. Review "rmdir."

Chapter 6: Working with System Files and the Boot Process

Exam Objectives

- ✓ Identify the names and purposes of each of the files that are required to boot up a Windows 2000 or Windows XP computer.
- ✓ Identify the locations of the Windows 2000 or Windows XP boot files.

The process of starting a computer has long been referred to as *booting*. Before you can use your computer, you need to be able to boot it to a point where the operating system (OS) is functional. Otherwise, your computer is like a safe without a known combination. This chapter will help you get that “safe” open by examining the boot process.

The boot process encompasses a series of steps, from the application of power to the loading of the OS shell. This chapter will review the hardware *POST (Power-On Self-Test)* process and will concentrate on the OS portion of the overall process. This chapter will also introduce you to the standard Windows XP boot process and the files that are required, and also how to correct or deal with boot problems related to the boot files. The Windows 9x/MS-DOS boot process will also be covered in this chapter, as you will likely use the MS-DOS boot process when creating your own boot disks and you may still encounter a number of Windows 9x computers in the field (although the number is falling rapidly). The Windows 9x/MS-DOS boot process is covered together, since MS-DOS still performs the role of boot loader for Windows 9x, so the initial process is the same for both. A solid understanding of the Windows XP boot files will prepare you for the exam, and a general knowledge of the other files will prepare you for working with these systems in the field. After covering the MS-DOS/Windows 9x boot process, you will look at memory management; good memory management can only help you if you need to revert to using an MS-DOS boot disk.

If, as an A+ Certified Professional, you are faced with troubleshooting the boot process on a computer for a user, it is required that you understand the boot sequence. If you don't know what the normal boot process is, then you are going to be at a deficit attempting to troubleshoot issues with it. As a CompTIA A+ Certified Professional, you will often have to bring your knowledge of the core operating system files to bear on your user's computer conundrums. Good knowledge of the MS-DOS boot files lets you work with and manage boot disks to help you recover damaged computer systems.



If you are unfamiliar with boot processes in general, focus on the entire chapter, but the exam will focus on the Windows XP boot process. The exam may include references to MS-DOS memory structures, the `config.sys` file, and the `autoexec.bat` file, and although these three items are not part of the Windows XP boot process, they are part of supporting MS-DOS applications running within a Windows XP *NTVDM* (*NT Virtual DOS Machine*), which you can review in Book VI, Chapter 2.

Power-On Self-Test (POST) Process

The *Power-On Self-Test (POST)* process starts when power is applied to the system. Electrical current makes its way from the power lead on the motherboard to the ROM-BIOS chips, and when the current is received by the BIOS chips they immediately begin executing their programs. One of the first checks is the memory (both a count and integrity test). After the memory check, the POST process moves on to find out what ports or I/O devices exist on the system. If the system is equipped with a *PnP BIOS* (*Plug and Play BIOS*), as most new systems are, then the BIOS-level PnP configuration takes place. The next thing that happens is a search for bootable disk devices. The order of this search is defined by the settings stored in CMOS memory, but is often: a: (floppy drive), c: (first partition on the first bootable hard drive), and then CD-ROM.

For each device in the list of potential bootable devices, the partition table is checked for the active partition. Floppy disks and the CD-ROM will only check the first partition. For this partition, the first sector is read and checked for a boot loader. For MS-DOS and Windows 9x the boot loader is `io.sys`, and for Windows 2000 and Windows XP the boot loader is `ntldr`. When this file is located, it is executed. If it was not found on the first potential bootable device, then the second and third devices are checked before reporting a boot failure.

Standard Boot Process for Windows XP

Windows XP—based OSeS have their roots in IBM's OS/2. It is partly due to this history that their boot process is very different from Windows 9x. Unlike Windows 9x, there is no real-mode boot component to the OS; Windows XP is a pure 32-bit OS. So the boot processes are similar, only until the boot loader is located at the end of the POST process.

The *boot sector* is created when the disk is formatted, and it contains a small program that has a mini file-system driver to read FAT, FAT32, and NTFS partitions. This program then looks for the real boot loader, which is `ntldr`.



Due to ARC naming conventions, Microsoft refers to the drive that has the boot sector on it as the *System Partition*, and to the partition that has the `windows` directory on it as the *Boot Partition*. To help keep these terms straight, remember that the OS does not really “boot” until `ntoskrnl.exe` is launched from the `windows` directory. The `windows` directory for a Windows 2000 computer is `winnt`.

Windows 2000 and Windows XP use many of the same files as MS-DOS and Windows 9x. However, many files are specific to Windows 2000 and Windows XP. The following sections provide an overview of Windows 2000– and Windows XP–specific files.

ntldr

`ntldr` is the *boot loader* for Windows XP. Its job is to coordinate the loading of the rest of the OS. `ntldr` is located on the root of your system partition, and if it is corrupted, it can easily be replaced from any other working copy of Windows XP. `ntldr` switches the memory model that is used on the system to a *flat memory model*, treating all memory on the system as one contiguous block. If your computer requires the `ntbootdd.sys` file, which is a SCSI controller driver, then it is loaded by `ntldr` so the rest of the boot process can access the boot drive. `ntldr` then reads the `boot.ini` file, if it exists, and displays the list of possible OSes that can be booted. If you want know what happens when the `boot.ini` file is missing, read through Book VII, Chapter 1.

After choosing any version of a Windows XP–based OS, `ntdetect.com` is called. `ntdetect.com` performs a hardware detection, scanning all hardware ports, processor make, model, and description, and the amount of RAM on the system. Once this information has been collected, it is returned to `ntldr` and will eventually make up the `HKEY_LOCAL_MACHINE\HARDWARE` key of the Registry.

The last step that is performed by `ntldr` is to launch `ntoskrnl.exe`. To launch `ntoskrnl.exe`, `ntldr` goes to the `system32` subdirectory of the directory that is listed in the `boot.ini` file.



When formatting a floppy disk using Windows XP, the boot sector is set to look for `ntldr`. If you leave a disk in your computer when it is being rebooted, you will see this message:

```
NTLDR is Missing
Press any key to restart.
```

For disks formatted with Windows 9x, this message will appear:

```
Invalid system disk
Replace the disk, and then press any key
```

ntbootdd.sys

If you boot your computer from an IDE/ATA drive or a SCSI drive from a SCSI controller that supports its own BIOS, you won't see the `ntbootdd.sys` file on your computer. However, if you are loading Windows XP from a SCSI drive on a controller that does not have its own BIOS, `ntbootdd.sys` appears on the root of the drive that has `ntldr` on it. `ntbootdd.sys` is the SCSI driver for your SCSI controller, but it has been renamed from the driver's actual name to `ntbootdd.sys`.

boot.ini

The `boot.ini` file is a text file on the root of your system partition that lists the OSes that are available to boot. It contains the default timeout value for the boot menu to be displayed, and where `ntldr` can find each copy of the OS. Upon looking at the `boot.ini` for the first time, you may be confused by the strange notation used to denote locations (this notation is called an ARC pathname and it is discussed in the next section).

Listing 6-1 shows what a sample `boot.ini` file would look like.

Listing 6-1: A Sample boot.ini File

```
[boot loader]
timeout=10
default= multi(0)disk(0)rdisk(0)partition(3)\WINDOWS
[operating systems]
C:\="Microsoft Windows 98"
multi(0)disk(0)rdisk(0)partition(3)\WINDOWS="Windows XP"
multi(0)disk(0)rdisk(0)partition(4)\WINNT="Windows 2000"
multi(0)disk(0)rdisk(0)partition(4)\WINNT="Win 2K Error"/SOS
multi(0)disk(0)rdisk(0)partition(2)\WINNT="Windows NT 4.0"
```

The first section in the `boot.ini` file is the `[boot loader]` section. It lists the timeout value, the number of seconds to display the boot menu, and the default OS to boot. The default OS is located in the OS listing in the file. The first OS in the list with a matching path is loaded by `ntldr`.

In this code listing, the `[operating systems]` section lists all of the OSes that `boot.ini` knows about on the system. This section would be built and added to as you install multiple copies of Windows XP—based OSes on a system. (You may have multiple versions of an OS installed if this is a testing or development system.) The sample `boot.ini` file above also has two entries that refer to the same path (`partition 4`). The difference between them is the application of the `SOS` switch at the end of the second line. The description strings that are used in this section are the display text for the boot menu, but they have no effect on the boot process of the OS.

If you have only one copy of Windows 2000 or Windows XP installed, then you won't see the boot menu at all, since there is only one choice of operating system to load.

ARC pathnames

Advanced RISC Computing Specification (ARC) is a rigid set of standards that has been presented by the *ACE (Advanced Computing Environment)* initiative. This initiative has been sponsored by major vendors in the industry, and one of the standards that arose was a naming convention to refer to disk partitions. To understand the parts of this name, refer to Table 6-1.

Table 6-1 ARC Path Components

<i>Type</i>	<i>Class</i>	<i>Description</i>
SCSI	Adapter	This is the ordinal number of the SCSI controller that is located in the system. The ordinal number refers to the order in which the controller was located. The hardware is scanned according to the buses that exist on the motherboard, and each bus is scanned starting with slot number 1. If controllers were located in slots 1 and 3, then the controller in slot 1 would be SCSI(0) and the controller in slot 3 would be SCSI(1). SCSI is only used in cases where the SCSI controller either does not have an on-board BIOS, or is disabled.
multi	Adapter	The ordinal number of the Multifunction controller in the system. Multifunction is used for all devices that do not use the previous listing for SCSI. This includes IDE controllers and SCSI controllers with the BIOS enabled.
Signature	Non-Classed	This notation can be used in place of SCSI or Multi, in order to help conform to Plug and Play specifications. Each drive that is identified by the Windows XP OS has a Signature written to it. The signature notation looks like <code>Signature(8765bfa4)</code> , and tells <code>ntldr</code> to look for a drive that has that signature and load the OS from there, regardless of which controller it happens to be found on.
disk	Controller	The SCSI ID number that has been assigned to the SCSI drive on the system. This is set to 0 when using the Multi (multifunction) adapter.
rdisk	Peripheral	The rigid disk number, referring to the physical location on the controller rather than the logical ID. The SCSI ID number refers to a logical ID.
partition	Block Device	The partition number for the partition that the OS will be found on. Partition(0) would refer to a drive with no partitions. Since Windows XP requires a partitioned drive to store files, the partition number will always be 1 or greater.

In the code from Listing 1-1 in the previous section, `multi(0)` referred to the first non-SCSI controller on the system; `disk(0)` did not refer to anything, as it would be a SCSI ID of a disk; `rdisk(0)` referred to the first disk on the controller; and `partition(4)` would be the fourth partition on the disk. Windows XP uses this specification since the assignment of drive letters to partitions is flexible, but partition locations are very rigid, thus preventing errors.

ntdetect.com

The next step after having chosen Windows XP from the `boot.ini` menu is to have `ntdetect.com` run. `ntdetect.com`'s only job is to find out what hardware is present on the system. This detection process is similar to what happens during the POST process at the hardware level. `ntdetect.com` checks for the following components:

- ◆ Bus/adaptor type
- ◆ Communication ports
- ◆ Computer ID
- ◆ Floating-point coprocessor
- ◆ Floppy disks
- ◆ Keyboard
- ◆ Mouse/pointing device
- ◆ Parallel ports
- ◆ SCSI adapters
- ◆ Video adapters

This information creates a hardware tree that is passed back to `ntldr` and eventually given to `ntoskrnl.exe`, which places it in the Registry.

ntoskrnl.exe

The main goal of the boot process is to get the operating system kernel loaded and functioning. The computer has already given you a choice of OSes, inventoried the hardware, and is now ready to actually start loading the OS into memory. `ntoskrnl.exe` represents the first and most important step in this process. The OS kernel for Windows XP is responsible for all thread level scheduling on the system. It plays a major control role, managing all of the other components on the system. Without it, there would be anarchy in the OS.

ntldr proceeds to the path that is specified in boot.ini to locate ntoskrnl.exe in the system32 folder. If ntldr locates the kernel, it proceeds to execute it. ntldr will generate a missing kernel error message if it fails to locate the kernel. Startup error messages are covered in Book VII, Chapter 1. Once the kernel is running, ntldr passes control of the system over to it. There are several steps to the kernel load, starting with loading devices, and then moving on to loading any system services. Once the services are running, it loads the default shell application and user profile.

The default user profile is used to run the user logon process. At this point, the logon screen will tell you to “Press Ctrl+Alt+Delete to begin.” After providing logon credentials, that user session is discarded, and a new one is started up for the new user.

The device load process

All the devices that are to be loaded during the system startup are listed in the Registry. The Registry includes information about each device in:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\`<device>`

In this location, you will find several values that describe how the device will start up. These values are listed in Table 6-2.

Table 6-2 Device and Service Settings in the Registry

<i>Value</i>	<i>Description</i>
Display Name	This string is used to display the name in areas of Windows XP, such as Device Manager or the Services administrative tool.
Error Control	This value configures how errors will be reported back to the OS. A value of 0 does not report any errors with the device back to the OS. A value of 1 reports errors normally. A value of 2 makes errors severe and will cause an automatic reboot of the computer to the “Last Known Good Configuration.” A value of 3 makes errors critical and will also cause an automatic reboot of the computer to the “Last Known Good Configuration.” If the “Last Known Good Configuration” is already being used, then severe errors will enable the computer to continue to boot, but critical errors will start the bug-check routine.
Group	Devices can be grouped together. This is done mostly for dependencies. If any device in a group fails, then dependent devices will not start up.
Image Path	This is the path and name of the actual driver file that is used for the device or service.
Start	This identifies when the device will start up. There are 5 start types: boot (0), system (1), automatic (2), manual (3), and disabled (4). Most devices have boot or system for a start value, but you may find a few that are set to automatic. Most services are set for either automatic or manual.

(continued)

Table 6-2 (continued)

<i>Value</i>	<i>Description</i>
Tag	A Tag ID is assigned to the service when it is installed, but is not actually used by the OS.
Type	Identifies that type of service or device. All devices should have a value of 1. Service types should be; 1 for Kernel device drivers, 2 for File System drivers, 4 for arguments for an adapter, 10 for single process Win32 applications that follow the Service Control Protocol, and 20 for Win32 Services that can share their process with other Win32 Services.

The service load process

All the services that are to be loaded during the system startup are listed in the Registry. The Registry includes information about each device in the following location:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\<service>
```

This is the same location that's used for devices. The biggest difference between devices and services is that devices map out to a physical piece of hardware, while services are only software. The values for services are listed in Table 6-2 in the previous section.

Loading the shell

After the processing of all the device drivers, the user's shell loads. The application that makes up the shell is actually defined by a `shell =` line in `system.ini`, the default being `explorer.exe`. If the current shell ever crashes and is removed from RAM, then `explorer.exe` will be loaded. Explorer checks the Registry to see what desktop components are supposed to be displayed and then checks the `HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\RunServices` key in the Registry to auto-start other applications. One of the services started at this point is the network service. When the network service is started up, you will be presented with a logon screen.

After loading the requested services, Explorer then executes any entries that it finds in the Registry in the `HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\RunOnce` key. Each entry under `runonce` is executed sequentially, waiting for each to finish before moving on to the next. When these are completed, Explorer then moves on to the `run` and `load` entries in `win.ini` to launch additional applications, followed by the `run` entries in the Registry, found in `HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Run`, and finally the Startup group from the Start menu, Start→Programs→Startup.



This process is almost identical for all versions of Windows that are based on Windows NT. Some versions may differ slightly, whereas the Windows 9x boot process is very different. For more on the Windows 9x boot process, proceed to the “*Understanding the Boot Process for MS-DOS and Windows 9x*” section, covered next in this chapter.

Understanding the Boot Process for MS-DOS and Windows 9x

This section will cover the boot process for Windows 9x, which has not changed much since the days of MS-DOS and Windows 3.11. It will discuss the difference in the process for Windows 95 and Windows 98, starting right after the Power-On Self-Test (POST) and continuing through to the loading of Windows 9x. When covering this material, you will see the options that relate to the version of MS-DOS that is included with Windows 9x, and not the last packaged version of MS-DOS, which was MS-DOS 6.22. On your job, you will most likely use the MS-DOS files from Windows 98, which has support FAT32 and is now the most popular version for people who are using MS-DOS boot disks.

io.sys

The MS-DOS/Windows 9x boot loader is `io.sys`. This file executes and proceeds to display the opening Windows splash screen, as well as launching the real-mode operating system components of Windows 9x. The Windows *splash screen* is the graphic with the cloud background and the animated band at the bottom of the screen that is displayed at the start of the boot processes. *Real-mode* is one of the operating modes of Windows 9x. It is used to process the initial boot of Windows 9x, to process boot files such as `autoexec.bat`, and to call `win.com`, which is the Windows GUI (*Graphical User Interface*). When in real-mode, the memory structures and restrictions that apply to MS-DOS apply to Windows 9x, which means that all components are limited to the first 640KB of memory. When you read “Managing Memory,” later in this chapter, you will learn about all of the limits that affect real-mode networking, and how to optimize your use of real-mode memory.

When loading the logo for display on the splash screen, Windows 9x first checks to see if a file (`c:\logo.sys`) exists. If it does exist, it displays the file; if it does not exist, then it uses a copy of the file that is embedded in `io.sys`.



If you do not like your startup splash screen, you can make your own by working with a bitmapped graphic that is 320x400 pixels in size, with 256 colors. Just save the graphic as `c:\logo.sys`. You can also create or modify `c:\windows\logow.sys` (“Windows is shutting down.”) and `c:\windows\`

`logos.sys` (“It is now safe to turn off your computer.”). Both of these files have the same dimensions as the `c:\logo.sys` graphic.

The `io.sys` file reads MS-DOS environment settings from `msdos.sys` and then `config.sys` before moving on to loading and passing control to `command.com`. `command.com` then calls on `autoexec.bat` and processes all of the commands in the batch file. At the end of processing `config.sys` (or if it does not exist), Windows 9x will enforce certain settings that are deemed necessary to Windows 9x, and it also ensures that certain required system-level device drivers (such as `himem.sys`) get loaded. If `autoexec.bat` does not exist, there are also certain elements that will be processed and enforced. The following section will describe `msdos.sys`, `config.sys`, and `autoexec.bat` in more detail. All of these files exist on the root of your bootable drive, usually `C:`.

You can’t configure `io.sys` directly, but when diagnosing boot problems, it is one of the files that you should check.



Remember that `io.sys` is OS version-specific. It has been written to load only one version of the OS, which means that there is a version to load MS-DOS 6.22, one unique version to load for Windows 95 Upgrade/Retail, Windows 95 OSR2, Windows 98, Windows 98SE (Second Edition), and Windows ME (Millennium Edition). Keep this in mind when you’re diagnosing boot problems.



The Windows 2000 and Windows XP boot processes do not use `io.sys`, but rather they use `ntldr` as the boot loader, which you will read about in “Standard Boot Process for Windows XP,” later in this chapter.

msdos.sys

The `msdos.sys` file is used to create the initial real-mode environment for Windows 9x. In MS-DOS 6.x and earlier versions this file was a binary file, but in Windows 9x this file is an ASCII file. If you open `msdos.sys` with any text editor (such as `edit.com` or Notepad) you will be able to see the current settings. You will also see a message stating that you should not delete the rows of x’s. Some anti-virus and other programs check the size of `msdos.sys` to ensure that it has not been tampered with. Oddly enough, these programs are really only testing for a minimum file size of 1,024 bytes. There are many settings that can be changed by editing the values found in this file.

The `msdos.sys` file has two major sections and resembles an `.ini` file, which is a standard Windows file that is used to store settings. The two sections are `[Paths]` and `[Options]`. The `Paths` section contains information about where Windows 9x is installed (Table 6-3 lists these settings).

<i>Value</i>	<i>Description</i>
WinDir	This is the Windows directory where most of the operating system files are run from. Default value is C:\Windows.
WinBootDir	This is usually the same location as WinDir. This option individually denotes the location of specific startup files. If WinDir is located on a drive other than C:, then WinBootDir may point to a location on the C: drive. Default value is C:\Windows.
HostWinBootDrv	This is the drive letter associated with the WinBootDir. Default value is C.
UninstallDir	This is the location of the files that should be replaced in the event of an uninstall of Windows 9x. The files are stored in w9xundo.dat and the original location information is stored in w9xundo.ini. Default value is C.

The Options section contains options about the boot process (Table 6-4 lists these settings). Most options only have two settings, 0 which is disabled and 1 which is enabled; if the option has additional settings, they are listed in the table.

<i>Option</i>	<i>Description</i>
AutoScan	Windows 95–OSR2 and Windows 98 support the automatic running of Scandisk. If this is set to 0, then scandisk.exe is disabled. A setting of 1 prompts the user to scan the system after an improper shutdown, while a setting of 2 will scan automatically. Default value is 1.
BootDelay	This sets the length of time (in seconds) that you have to press one of the boot keys when you see the text “Starting Windows 95 . . .” at boot. Default value is 2.
BootSafe	A setting of 1 forces a safe mode boot of Windows 9x. Default value is 0.
BootGUI	A setting of 0 will cause the boot process to stop after processing autoexec.bat. You will be able to launch Windows 9x by typing WIN. This is useful during troubleshooting the boot process. Default value is 1.
BootKeys	Enables (1) or disables (0) the ability to press the boot keys during boot. The boot keys include F8, F5, F6, and the Ctrl key. Default value is 1.
BootMenu	Automatically displays the Windows 9x boot menu if set to 1. Default value is 0.

(continued)

Table 6-4 (continued)

<i>Option</i>	<i>Description</i>
BootMenuDefault	Automatically pre-selects a choice in the boot menu. Default value is 1.
BootMenuDelay	Sets the timeout value for the boot menu. Default value is 30.
BootMulti	Enables the option in the boot to the previous version of MS-DOS when set to 1. This option should not be enabled on OEM versions of Windows 9x, as they do not have a previous version of MS-DOS. Default value is 1.
BootWarn	Enables the Safe-Mode warning message when set to 1. Default value is 1.
BootWin	A setting of <code>BootWin = 0</code> will cause the system to automatically boot to the previous version of MS-DOS. This option should not be enabled on OEM versions of Windows 9x. Default value is 1.
DoubleBuffer	Enables the double-buffering driver for SCSI controllers when set to 1. Default value is 0.
DBLSpace	Loads the Double Space driver (<code>dblSPACE.bin</code>) if it exists on the root of your drive, when this is set to 1. Default value is 1.
DRVSpace	Loads the Drive Space driver (<code>drvSPACE.bin</code>) if it exists on the root of your drive, when this is set to 1. Default value is 1.
LoadTop	By default, Windows 9x will load <code>command.com</code> and <code>drvSPACE.bin</code> to the top of the 640K memory range. A value of 0 causes these files to load at the bottom of the range. Some applications, such as the Novell Netware client, require that <code>command.com</code> be loaded low. Default value is 1.
Logo	A value of 0 will disable the animated logo. The animated logo can cause problems with some TSR (Terminate Stay Resident) programs or memory managers. Default value is 1.
Network	A value of 1 will load network support when entering Safe Mode. Default value is 0.



Windows 2000 and Windows XP boot processes do not use the `msdos.sys` file.

config.sys

The `config.sys` file has been around since early versions of MS-DOS. It stores modifications to the default settings that are set through `msdos.sys`. `config.sys` usually contains environment settings and device drivers load lines, for devices such as CD-ROM drives. If the file does not exist, then `io.sys` will apply system default settings to the environment, in addition to the settings that are found in `msdos.sys`. If the `config.sys` file does exist,

`io.sys` ensures that after processing the `config.sys` file, the settings are equal to or higher than the default settings. If `config.sys` settings do not meet the default settings, then `io.sys` will apply the default settings. (Your settings may have been changed from the default settings in order to support an application; the file's value is often increased for application support.) The default entries that will be enforced are found in Table 6-5.

Table 6-5	Default config.sys Values
<i>Entry</i>	<i>Description</i>
<code>buffers=23</code>	Buffers are used when some older MS-DOS—based applications attempt to make file I/O calls. They are not used by Windows 9x or its applications.
<code>dos=High, Auto</code>	Loads some of the <code>command.com</code> and <code>msdos.sys</code> files into the High Memory Area.
<code>files=30</code>	Like <code>buffers</code> , this setting is only used by older applications. <code>files</code> sets the number of files that the application or <code>command.com</code> environment can have open simultaneously.
<code>himem.sys</code>	High Memory Manager. This file gives Windows 9x access to memory above 1MB. In a round-about way, it also includes all memory above 640K.
<code>ifshlp.sys</code>	Installable File System Driver Helper. This driver is required to allow access to both network file systems and VFAT (the local file system).
<code>lastdrive=Z</code>	This gives access to drive letters up to Z:. It is only required for older applications.
<code>setver.exe</code>	Emulates different versions of MS-DOS for compatibility with some applications.
<code>shell=command.com /p</code>	Sets the command shell to be permanent.
<code>stacks=9,256</code>	Like <code>buffers</code> , this is used for compatibility with some older MS-DOS—based applications.

In Chapter 1 of this minibook, you are told that Windows NT can be traced back to a joint venture with IBM, and originally ran on IBM's OS/2 (Operating System 2). Since that time, Microsoft has provided some support for OS/2 applications in the Windows NT OS. This support was removed with the release of Windows XP. Windows 2000 and Windows XP do not generally use the `config.sys` file. If the file is edited and saved as an OS/2 text file, Windows 2000 will extract any OS/2 configuration information. This information will be written to the Registry and used for running OS/2 applications from within the Windows 2000 OS. This process happens only with Windows NT and Windows 2000 and not with MS-DOS or Windows 9x.

command.com

The `command.com` file is the command interpreter for Windows 9x. The job of the *command interpreter* is to execute non-graphical applications for Windows 9x. It is located in several places on your hard drive: `C:\command.com` and `C:\windows\command.com`, and possibly either `C:\windows\command\command.com` or `C:\windows\command\ebd\command.com`. These additional files provide a small amount of fault tolerance to this file. Windows 9x usually uses the copy located in the Windows folder rather than the copy located on the root of the `C:` drive. The copy on the root of the drive is only used during the first stage of the boot process, the real-mode portion. `command.com` is loaded by `io.sys` and is responsible for the execution of all commands prior to the loading and initialization of the Windows kernel with the execution of `win.com`.

If something happens during the boot process that prevents `win.com` from executing, or if you are using a MS-DOS boot disk and you do not load the Windows 9x GUI (Graphical User Interface) — you are faced with a command prompt provided by `command.com`. If you are unfamiliar with the command prompt, you may ask yourself, “Now what?” The command prompt, which is composed of a `C:\>` and a hypnotically blinking underscore, looks back at you unyieldingly. Unfortunately, it won’t share its secrets unless you ask nicely!



If you have a full version of MS-DOS installed on your computer, then you can get help for all commands by typing `help` or `help` with a specific command by typing `help` and the name of the command, like this:

```
help xcopy
```

Rather than using the detailed `help` command, try typing the name of the command followed by the `/?` switch. This option will usually give you an abridged version of the help information.



If you have opened a command prompt from within Windows 2000 or Windows XP, you will find help for many MS-DOS commands in the Windows Help file (Start⇨Help or Start⇨Help and Support).

Asking for help from the command prompt implies that you know what to ask, or at least that you know the name of the command that you need help with. In the case of the GUI, you are given hints and a fairly intuitive way of moving around and accessing information. However, basic functionality at the command prompt requires only that you know a small handful of commands. These commands, covered in the previous chapter, include navigation commands such as `dir`, `cd`, and `md`; movement commands such as `ren`, `copy`, and `xcopy.exe`; disk-level commands such as `format.com`; and text manipulation commands such as `type` and `edit.com`.

All of these commands happen at the command prompt. You will often have to revert to the command prompt within Windows 2000 or Windows XP to execute command-line programs. In the case of Windows 9x, you may also have to resort to the command prompt when you are unable to load the GUI due to errors in the system configuration.

autoexec.bat

After `command.com` loads and initializes, it is configured to look for and process `autoexec.bat`. `autoexec.bat` is a batch file. *Batch files* simply store a list of commands that are executed in order, and `autoexec.bat` is configured to be executed at boot. In addition to executing programs, `autoexec.bat` is used to set environment variables and load *TSRs* (*Terminate Stay Resident programs*).



TSRs are applications that load and run but do not normally occupy space in the user interface. A good example is `doskey.com`. When you run `doskey.com` from a command prompt, it doesn't appear to do anything, but it actually does: It runs, loads itself into memory, and then returns the command prompt back to you. This return of control is different from a command such as `xcopy.exe`, which controls the command prompt while it runs. It then completes its task and exits (terminates). When it terminates, it removes itself from memory and is no longer running.

`autoexec.bat` is used to set environment variables, such as the location of the temporary directory and the search path. Windows 9x will enforce certain variables and settings prior to executing `autoexec.bat`, which allows you to change, append, or replace any of these values. The default values for the Windows 9x `autoexec.bat` are listed in Table 6-6. Since the temporary directory is set to the Windows directory, many people change at least this one path to point to a different location, so that temporary files do not fill the Windows directory. If you want to add to the default system path, you can do so with a statement such as:

```
SET PATH = %PATH%;C:\MYAPPS;
```

This will add to the existing path (`%path%`), rather than overwrite it.

If `win.com` is not called during `autoexec.bat`, then `command.com` will process the call to `win.com`.

Table 6-6 Default `autoexec.bat` Values

<i>Variable</i>	<i>Value</i>
Tmp	C:\windows
Temp	C:\windows

(continued)

Table 6-6 (continued)

Path	C:\windows;C:\windows\command
Prompt	\$p\$g
Comspec	C:\windows\command.com



As with the `config.sys` file, MS-DOS and Windows 9x can function without `autoexec.bat` — though if the `autoexec.bat` file is missing under MS-DOS, you are prompted to confirm the date at every boot-up.



Windows 2000 and XP will read the `autoexec.bat` file, but they ignore lines that execute programs or load TSRs. The only items or lines that are used out of the `autoexec.bat` file are statements that set environment variables, such as the system path.

win.com and vmm32

The `win.com` file can be called through `autoexec.bat` or from the command line. When called through `autoexec.bat`, you will be able to pass switches to `win.com` in order to disable certain features. `win.com` immediately turns control over to `vmm32` (*Virtual Memory Manager*), which proceeds to load the Windows 9x *Graphical User Interface (GUI)*. `vmm32` starts by changing the memory model that is being used, and gives the OS direct access to all physical memory above 1MB, while still supporting any real-mode components that were loaded below 1MB. `vmm32` then scans the system Registry and attempts to load any devices that are listed in there. One driver that will be listed in the Registry is `vmm32.vxd`.

During setup, all required drivers are merged into a single file named `vmm32.vxd`. If additional drivers are required or loaded after setup, then they are stored in the `C:\windows\system\vmm32` directory. If setup is re-executed later, then these files are merged into `vmm32.vxd`. Due to this handling of `vmm32.vxd`, it is a file that is specific to each workstation.

After `vmm32` loads `vmm32.vxd` and the other drivers listed in the Registry, it then attempts to load any devices that are listed in the `[386enh]` section of `system.ini`. Once all of the hardware devices have been identified, `vmm32` places the processor into Protected-Mode and loads the OS Kernel (`krnl386.exe` and `kernel32.dll`), followed by `gdi.exe`, `gdi32.dll`, `user.exe`, and `user32.dll`. Then system resources and fonts are initialized. The last step is to check the `win.ini` file to see if there are any additional settings that should be enforced on the system, and what the shell application is supposed to be. Table 6-7 examines each of these components.

Table 6-7 Components Used When Loading Windows 9x	
Component	Description
Registry	A registry is a place to record information. For example, at a university you can find out what students are attending which classes by consulting the Registrar's Office. The same is true with Windows 9x and Windows XP-based OSes. The Registry is used to store information that is used by other components on the system. In Windows 9x, the Registry is composed of two OS files: <code>user.dat</code> and <code>system.dat</code> . It is used by <code>vm32</code> to identify devices that are supposed to be loaded or initialized on the system. Settings for each device are also located in the Registry. Great care must be taken when editing the Registry.
<code>system.ini</code>	<code>system.ini</code> contains loading information for all real-mode devices. In most cases, you will find that this includes your mouse and video drivers.
<code>kernel32.dll</code>	<code>kernel32.dll</code> is the core set of code that makes up the Windows 9x operating system. It takes over from <code>vm32</code> for managing the system. It schedules and manages all other processes or applications that are running on the system.
<code>krnl386.exe</code>	This component exists on the system for backward compatibility with older Windows 3.x programs that want to pass calls to <code>krnl386.exe</code> . Any calls that are passed to <code>krnl386.exe</code> are redirected to <code>kernel32.dll</code> .
<code>gdi.exe</code>	This is the 16-bit component that is responsible for handling the Graphic Device Interface. If there are graphics (such as windows) that must be presented on the screen or on a printer, then this is the component that is responsible for it.
<code>gdi32.dll</code>	This is the 32-bit version of <code>gdi.exe</code> . Functionality is not duplicated in these two GDI components; rather, some features are implemented in one or the other. Both components accept all of the component calls, but may pass the request to the other component. A program is able to call on a component using a 32-bit GDI call. If that feature is actually in <code>gdi.exe</code> , then <code>gdi32.dll</code> will pass the call to <code>gdi.exe</code> and process the response. This allows a programmer to program to one interface, regardless of where Microsoft actually stored the components.
<code>user.exe</code>	This is the 16-bit component that is responsible for user input. Most of the functionality of the user interface components is located in <code>user.exe</code> .
<code>user32.dll</code>	This is the 32-bit version of <code>user.exe</code> . Since most of the functionality has been implemented in <code>user.exe</code> , <code>user32.dll</code> passes most of its calls over to <code>user.exe</code> .
Resources and fonts	System resources are reserved for most of the main system components. Fonts actually make up an integral part of the OS, so they are assigned system resources at this point, along with some other components.

(continued)

Table 6-7 (continued)

win.ini values	win.ini stores information on a number of shell-related settings. Some shell settings are stored in the Registry, but for compatibility with older applications, some of these settings are also found in win.ini.
----------------	--



For information on editing the Registry, see the `regedit.exe` section in Book VI, Chapter 4.

Loading the shell

At this part in the boot process, Windows 9x will load the shell application. This part of the boot process is virtually identical to the loading of the shell for Windows XP, so you can refer back to the end of the Windows XP boot process to see what happens when the shell is loaded.

Managing Memory

If you are lucky enough never to have created a boot disk for your computer, then you have never had to get into the down-and-dirty world of memory management. When you are attempting to load device drivers and applications in the MS-DOS environment, you have a very limited amount of the memory to work with — 640KB — and if you are not careful, you can quickly run out of space; so you need to manage what and how device drivers and applications get loaded into memory. If you have to support MS-DOS-based applications within the Windows environment, or if you need to create custom boot disks that will load your network drivers, you need to pay attention to this section. Memory management skills are becoming a lost art in the world of Windows.

Conventional memory

When MS-DOS was first developed, the world was just starting to imagine that computers would actually have 1MB of memory. The first IBM PC released in 1981 shipped with only 320K of memory. Within 18 months of its release, the IBM PC was followed up by a new version of MS-DOS and an unheard of 1MB of memory. It was seriously believed (shortsightedly) that computers would not progress beyond the 1MB mark in the near future. A decision was made early on to slice this 1MB pie into a couple of pieces. This slice cut out the first 640K of memory and called it *conventional memory*.

Conventional memory was “conventional” because it was the place where normal applications would run. All applications run in the conventional memory space. Many TSRs and drivers also load into this space, each taking a piece of the space away from other applications. Because a finite amount of memory is available in this 640K area, care has to be taken to load the

required drivers while leaving space for applications to run. Even with Windows 9x, conventional memory will still have an effect on you, until you launch the `win.com` and `vmm32` changes the memory model that is being used.

The rest of the memory below the 1MB mark was to be used for loading device drivers and supporting different levels of video displays.

Expanded memory

The limit of 640K of conventional memory became troublesome for many people, and by 1984, several methods were available for expanding the memory on your computer. By adding memory on expansion cards, more memory could be made available to applications. Lotus, Intel, and Microsoft created the *LIM EMS (Lotus Intel Microsoft Expanded Memory Specification)* standard for the implementation of memory expansion. This is often just called EMS. Memory added through expansion cards was — very logically — called *expanded memory*. Although all of the system's “action” took place in conventional memory, the expanded memory gave applications an area where they could store information.

Expanded memory could be beneficial only if applications were designed to use it, and eventually, many applications grew to rely on the extra storage space provided by expanded memory. Expanded memory resulted in increased application performance. After the introduction of the 80286 processor (1984), preference was given to extended memory (see the following section), and expanded memory cards became a thing of the past. Nowadays, applications that require expanded memory use an emulator, `emmm386.exe`. For more on `emmm386.exe`, see the section “`emmm386.exe`,” later in this chapter.

Extended memory

Eventually, computer manufacturers devised a way to enable hardware to address more than 1MB of memory, and the amount of available memory then jumped to 2MB or 4MB. Unfortunately, the MS-DOS real-mode operating system was hard-coded to limit applications to 640K. In order to enable access to the additional memory, and at the same time allow Windows to use the additional memory, the idea of a memory manager was born. The *memory manager* would control access to the memory that was extended beyond the original 1MB memory chips.

Allowing a memory manager, such as `himem.sys`, to swap information in and out of extended memory gave applications access to a larger block of memory. All memory that was being accessed had to be below the 640K mark, but it could be swapped into the extended memory area at other times. With the release of protected-mode Windows 3.0 (1990), Windows was able to access extended memory directly.

High memory

When the driver (`himem.sys`) was created to enable access to extended memory, they took advantage of a processor design flaw, which started with the 286 processor. With MS-DOS 5.0, the `himem.sys` developers found out how to exploit this processor flaw to allow MS-DOS to access the first 64K of extended memory (from 1,024K to 1,088K). This area was referred to as the *High Memory Area (HMA)*.

MS-DOS is able to load a portion of its code into the High Memory Area, freeing up memory below 640KB for other drivers and applications to use. This is a standard technique used to allow additional applications to run in the MS-DOS environment. To give MS-DOS access to this area, you will need to add the line `dos = high` to your `config.sys`. Your `config.sys` file can be modified with any text editor. To then have device drivers load into the HMA, you will change your device load commands in `config.sys`. The following two lines show a standard device load command, and a load command which uses the HMA.

```
device=a:\ansi.sys  
devicehigh=a:\ansi.sys
```

Upper memory

Upper memory refers to a portion of the memory that exists between the 640K and 1MB marks. A large portion of this area was originally allocated for use by system devices, such as your video display. A Windows 9x computer can use this area to emulate expanded memory, to load drivers, or both. When you use this area for either of these purposes, you'll need to load the `himem.sys` and `emm386.exe` drivers because high memory is used to gain access to the upper memory area. Both of these drivers — and some of their options — are covered in the sections “`himem.sys`” and “`emm386.exe`,” later in this chapter.

Virtual memory

As improvements were made in the field of RAM, and as computers with more and more memory continued to ship, software developers created applications that used the new memory. To make the entire process of managing memory easier, Microsoft decided to implement virtual memory for the Windows operating systems. Virtual memory allows Windows to present a Virtual Machine (VM) that contains 4GB of memory to applications running in the Windows environment. It then used a *Virtual Memory Manager (VMM)* to control or manage the mapping of data between the virtual addresses used by the application and where the data was stored in physical memory. The VMM

was also able to move data that was not being actively used in RAM to a file on the disk. The swapping of memory data pages to and from the disk file lead to the file being named swap file in Windows 9x and paging file in Windows NT. The drawback in the system of swapping data shows up when an application wants to use data that is in the swap file, since it then has to wait for the data to be retrieved back into RAM before it can be accessed.

Access speeds of hard disks are measured in milliseconds (10^{-3}), while memory access is measured in nanoseconds (10^{-9}). It should not be hard to guess that this means that when data has to be retrieved from the swap file on a hard drive, the process is extremely slow relative to retrieving it directly from RAM.



The Virtual Memory Manager manages virtual memory addresses up to 4GB and the mapping of those addresses to a physical location, either in RAM or a hard drive. You should rely on using the paging file only when applications need a small amount of additional memory. Most operating systems implement virtual memory and allow it to use swap space on a hard drive to allow applications with high memory requirements to function, but greater performance will be achieved by adding more physical RAM to the system.

When an application needs to store information to memory, it passes the request to the VMM. VMM stores the information in RAM but may move the information to the swap file on the drive at a later time. The process for retrieving application data is illustrated in Figure 6-1; the process looks like this:

1. When the application requests information, the VMM checks to see whether the information is in RAM.
2. If the information is in RAM, the information is simply returned to the application, and the process is complete.
3. If the information isn't in RAM, VMM checks to see if there is enough space in RAM to retrieve the information from the swap file.
4. If there is enough space to retrieve the information, then the information is retrieved from the drive, stored into RAM, and passed on to the application, and the process is complete.
5. If there isn't enough space to retrieve the information, VMM checks for memory locations that have not been accessed recently and passes them from RAM down to the swap file.
6. When enough information is moved to the swap file to make room for the requested information, that information is moved into RAM and then returned to the application.

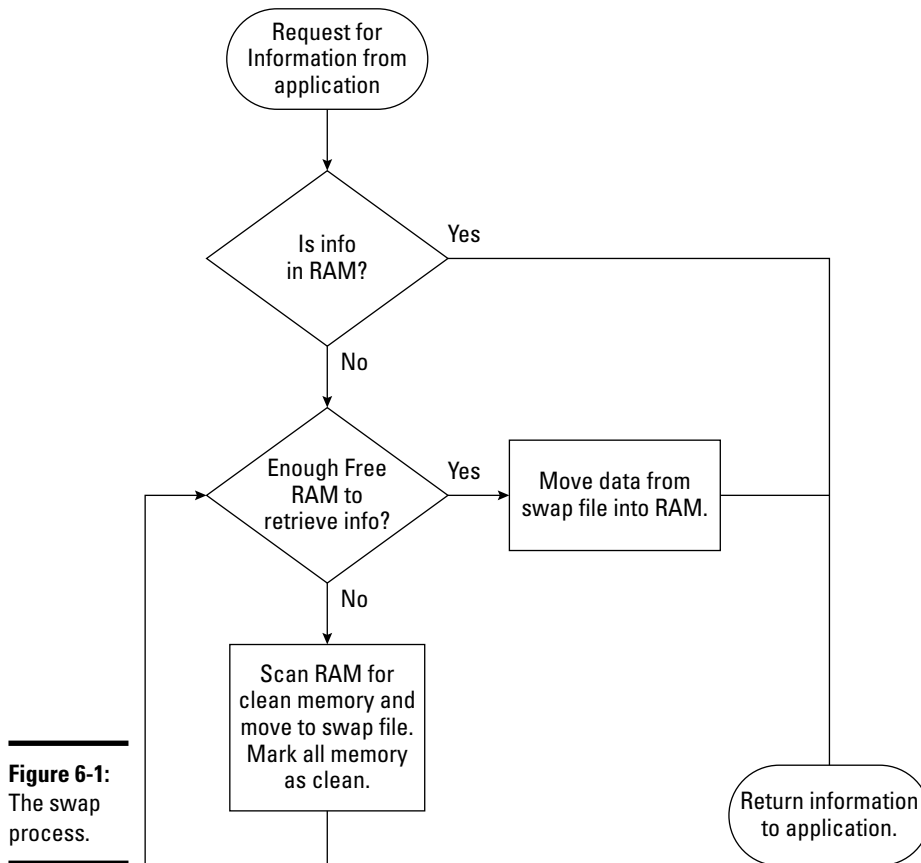


Figure 6-1:
The swap process.

A *clean* memory location in RAM is a location that has not been accessed since the last time the VMM marked it clean. If the memory location has been accessed with a read or write request, then this location is marked as *dirty*. When looking for memory data to move to the swap file, each location is checked; if it is clean it is moved to the hard drive, and if it is dirty, it is marked as dirty and left. If the first scan did not free enough RAM, then an immediate second search for movable memory data is required, at which time any memory data that is now dirty is data that was accessed since the first scan, mere milliseconds ago. This algorithm is called the *Least Recently Used (LRU)* algorithm, and it ensures that data that is actively used in RAM will stay in RAM.

himem.sys

`himem.sys` is a high memory manager that allows computers to access memory above the 1MB mark. It is loaded in the `config.sys` file with the following line:

```
device=c:\windows\himem.sys
```

An optional switch (`/testram:off`) can be used to bypass the memory test. This test is considered redundant by some people because RAM is tested during the POST (Power-On Self Test) process.

On Windows 9x systems, `himem.sys` is automatically loaded at the end of processing `config.sys` — that is, if it wasn't already loaded. The Windows 9x GUI requires `himem.sys`.

The addition of the line `dos=high` anywhere in your `config.sys` file causes a portion of `command.com` to be loaded into the HMA. Windows 9x systems automatically process `dos=high` if it is not present in the `config.sys` file.

emm386.exe

Windows does not require `emm386.exe`, but it can be used to optimize memory configurations or to support MS-DOS-based applications.

If you want to use `emm386.exe` to support your applications that require expanded memory, you can load `emm386.exe` in your `config.sys` file by editing `config.sys` with any text editor and adding the following line:

```
device=c:\windows\emm386.exe ram
```

Loading `emm386.exe` in this way creates a 64K page frame in the upper memory. This page frame comprises four 16K pages. `emm386.exe` uses extended memory to simulate expanded memory, but information that is being manipulated must reside below the 1MB mark of memory. To achieve this, information from extended memory is swapped into the page frame 16K at a time.

If you know that none of your applications require EMS, you can make additional space available as *Upper Memory Blocks (UMB)* by adding the following lines in `config.sys` using a text editor of your choice:

```
dos=umb  
device=c:\windows\emm386 noems
```

Not all driver files and TSRs can be loaded into UMBs, but for those that can, you can modify device lines in `config.sys` to `devicehigh`, as shown in this example:


```
devicehigh=c:\windows\command\ansi.sys
```

You can also modify TSR lines in `autoexec.bat` using a text editor and placing an `lh` (load high) at the beginning of the line, like this:

```
lh c:\windows\smartdrv.exe
```

You can check on your success from the command prompt by typing the `mem /c` command. This gives you a listing of what has loaded into conventional memory and what has loaded into upper memory. Any drivers or TSRs that loaded into UMBs will have reduced conventional memory requirements. The following is a sample of the output from the `mem` command.

Modules using memory below 1 MB:

Name	Total	Conventional	Upper Memory
SYSTEM	34,592 (34K)	10,656 (10K)	23,936 (23K)
HIMEM	1,168 (1K)	1,168 (1K)	0 (0K)
EMM386	4,320 (4K)	4,320 (4K)	0 (0K)
COMMAND	8,224 (8K)	1,056 (1K)	7,168 (7K)
RAMDRIVE	1,456 (1K)	0 (0K)	1,456 (1K)
ANSI	4,320 (4K)	0 (0K)	4,320 (4K)
DOSKEY	4,688 (5K)	0 (0K)	4,688 (5K)
SMARTDRV	32,192 (31K)	0 (0K)	32,192 (31K)
Free	722,624 (706K)	638,080 (623K)	84,544 (83K)

Memory Summary:

Type of Memory	Total	Used	Free
Conventional	655,360	17,280	638,080
Upper	158,304	73,760	84,544
Reserved	393,216	393,216	0
Extended (XMS)	82,679,200	8,717,728	73,961,472
Total memory	83,886,080	9,201,984	74,684,096
Total under 1 MB	813,664	91,040	722,624
Largest executable program size		638,064 (623K)	
Largest free upper memory block		84,256 (82K)	

MS-DOS is resident in the high memory area.

You should note that drivers such as `ramdrive.sys` and `ansi.sys` have loaded into upper memory, and TSRs such as `doskey.exe` have also loaded into upper memory.

By using the `include` switch, you can include areas of upper memory that are reserved for other components. For example, `i=B000-B7FF` would include the area that is reserved for monochrome monitors, and `i=E000-EFFF` would include an area that is reserved for IBM PS/2 mice. An example of the full line `emm386.exe` line in `config.sys` would be:

```
device=a:\emm386.exe i=E000-EFFF i=B000-B7FF noems
```



To allow you to see the effect of using `himem.sys` and `emm386.exe` to optimize conventional memory for applications, Lab 6-1 will walk you through the process of memory optimization one step at a time. Lab 6-1 can be found in the `Labs.pdf` file in the `Author` directory of the CD-ROM.

Examining Other Boot Process Files

This section will take a look at some additional files that are related to the boot process. These will include `smartdrv.exe`, the MS-DOS disk caching utility which improves disk access times; `win.ini` and `system.ini` which store configuration information for Windows, primarily Windows 9x, but they are also used by Windows 2000 and Windows XP; finally you will see the `sysedit.exe` utility.

smartdrv.exe

Caching within computers has been implemented at different levels. *Smart drive* (`smartdrv.exe`) represents one of those levels. `smartdrv.exe` is a disk-caching program, which means that it reads additional data from the drive in anticipation of your program's request. Anticipating reads is the *read caching* feature of `smartdrv.exe`. `smartdrv.exe` also performs *write caching*, which is when it collects the write requests from several seconds and then commits all of the data to the disk at once. *Disk caching* improves drive I/O performance by reading and writing to the drive in larger units, which requires less drive head movement than performing reads and writes in smaller units.

`smartdrv.exe` is usually loaded only when creating a boot disk or installing Windows 9x. Windows 9x has its own built-in disk caching software. If you have some older applications that won't work with Windows 9x's disk caching, you can disable it and load `smartdrv.exe`. To disable the Windows 9x disk caching, right-click `My Computer` → `Performance` → `File System` → `Troubleshooting` → `Disable write-behind caching for all drives` → `OK` → `OK`.

`smartdrv.exe` can be loaded through either `config.sys` or `autoexec.bat`. If you have an old disk controller that will not work with memory that is provided by `emm386.exe`, then you will have to load `smartdrv.exe` through `config.sys` with the `double_buffer` option. Double buffering will only be required on a few very old SCSI hard drives, so you should not expect to encounter a time when you will need this option. To load `smartdrv.exe` with this option, add the following line to your `config.sys`, by editing the file with a text editor:

```
device=c:\dos\smartdrv.exe /double_buffer
```

The `smartdrv.exe` line must appear before the line that loads `emm386.exe` if you are using the `double_buffer` option.

Most of the time, you will load `smartdrv.exe` from the command prompt or through `autoexec.bat`. Table 6-8 contains the main options that you will use with `smartdrv.exe`.

Table 6-8 smartdrv.exe Switches	
Option	Description
<i>drive_letter</i> + or -	Lets you enable or disable caching on specific drives by using this command: <code>smartdrv.exe a- c+ d-</code> When you load <code>smartdrv.exe</code> , it automatically enables read and write caching on all hard drives, read caching on floppy drives and CD-ROM drives, and it ignores network drives and flash memory-card drives.
Initial Cache Size	If you specify a size in kilobytes after the <code>smartdrv.exe</code> command, then it will be treated as the initial cache size. This tells <code>smartdrv.exe</code> how big to make the cache. (Table 6-9 describes how the default cache sizes are set.) This value can be set as high as 16,383K. In general, the larger the cache, the less often you will have to read from the drive. To set the cache size to 2MB, you would use the command <code>smartdrv.exe 2048</code> .
Windows Cache Size	Specifying only one size in kilobytes after the <code>smartdrv.exe</code> command sets the initial cache size, and the Windows cache size is still based on the values in Table 6-9. If you want to control the Windows cache size as well as the initial cache size, you provide two values after the <code>smartdrv.exe</code> command. To set the initial and Windows caches to 4MB, you would use this command: <code>smartdrv.exe 4096 4096</code> . The initial cache is always the first number and it has to be as large or larger than the Windows cache. If you set the Windows cache to a value that is higher than the initial cache, <code>smartdrv.exe</code> uses the Windows cache size for the initial cache size.
/X	Disables write caching on all drives. It takes several seconds after your application tells <code>smartdrv.exe</code> to write to the drive before the information is actually written to the drive. If the power is interrupted during this time, the data will be lost. If you are concerned about the loss of data, then you will want to disable write caching.
/C	This switch is used from the command prompt to instruct <code>smartdrv.exe</code> to write all outstanding information to the drive.
/V	Verbose mode makes <code>smartdrv.exe</code> display status and error messages when it loads.
/Q	Quiet mode makes <code>smartdrv.exe</code> hide status messages when it loads. It still displays error messages.
/S	This switch is used from the command prompt to display the status of <code>smartdrv.exe</code> .



When using MS-DOS, you should make sure that your computer has finished writing information to the drive, prior to turning off your computer; you can ensure that `smartdrv.exe` has committed all data to the drive by typing `smartdrv.exe /c` at the command prompt. `smartdrv.exe` will automatically commit data to your drive if you reboot by pressing Ctrl+Alt+Delete.

The initial cache size and Windows cache sizes can be set when loading `smartdrv.exe` but will default to the values listed in Table 6-9. The default values are based on the amount of extended memory on the computer.

<i>Amount of Extended Memory</i>	<i>Initial Cache Size</i>	<i>Windows Cache Size</i>
0–1MB	All extended memory	0
1–2MB	1MB	256K
2–4MB	1MB	512K
4–6MB	2MB	1MB
6MB and up	2MB	2MB

system.ini

Before Windows 95, all configuration settings for applications were stored in `.ini` files. All `.ini` files have a similar structure, which is represented in the following code section:

```
[section one]
setting one = "my mystery word"
setting two = 42

[section two]
setting one = "another mystery word"
setting two = 75
```

The `system.ini` file contains settings that are used specifically by Windows. The type information that is found in this file includes device settings and settings that are used when loading the Windows kernel. The `system.ini` file isn't used much these days, compared to its importance in Windows 3.x, but it still finds use when you need to load 16-bit MS-DOS real-mode drivers. If you check your `system.ini` file, you should see settings for your mouse and video. Here is a sample of what to expect:

```
[386Enh]
device=*vshare
device=*dynapage
device=*vcd
device=*vpd
device=*int13
mouse=*vmouse, msmouse.vxd
woafont=dosapp.fon
COM1Fifo=0
keyboard=*vkd
```

```
display=*vdd,*vflatd

[boot]
system.driv=system.driv
drivers=mmsystem.dll power.driv
user.exe=user.exe
gdi.exe=gdi.exe
sound.driv=mmsound.driv
dibeng.driv=dibeng.dll
comm.driv=comm.driv
shell=Explorer.exe
keyboard.driv=keyboard.driv
fonts.fon=vgasys.fon
fixedfon.fon=vgafix.fon
oemfonts.fon=vgaoem.fon
386Grabber=vgafull.3gr
display.driv=pnpdrv.r.driv
mouse.driv=mouse.driv
*DisplayFallback=0
SCRNSAVE.EXE=
```

In this example `system.ini` file, you can see these settings:

- ◆ In the `[386Enh]` section,
 - File and printer sharing is initialized with `*vshare`.
 - Real-mode support for the mouse and video display is loaded with `*vmouse`, `msmouse.vxd`, `*vdd`, and `*vflatd`.
- ◆ In the `[boot]` section,
 - Power management support is loaded with `power.driv`.
 - The sound card is initialized with `mmsound.driv`.
 - Serial communication is initialized with `comm.driv`.
 - The system shell application is set with `shell=Explorer.exe`.

If the `system.ini` file is lost or corrupted, you will likely have problems with all the devices that rely on it. Because Windows 9x stores some video configuration in this file, you tend to see problems quickly. If the `system.ini` file is missing, all versions of Windows create a new one and start loading the settings back in. This file is not used by the MS-DOS boot process.

win.ini

Like `system.ini`, `win.ini` was heavily used by both Windows 3.x and its applications. Today, it is used primarily for backward compatibility with

older applications, although it still stores some settings. Here is a sample of some of the information found in a `win.ini` file:

```
[windows]
run=Qtstub.exe
NullPort=None
ScreenSaveActive=1
UninstallPath=C:\
device=EPSON Stylus COLOR 600,EPS600,LPT1:

[Desktop]
Wallpaper=C:\WINDOWS\APPLIC~1\MICROS~1\INTERN~1\INTERN~1.BMP
TileWallpaper=0
WallpaperStyle=0
Pattern=(None)

[Extensions]
txt=notepad.exe ^.txt
bmp=C:\Progra~1\Access~1\mspaint.exe ^.bmp
vdo=VDOLIV32.EXE ^.vdo
gra=C:\PROGRA~1\COMMON~1\MICROS~1\MSGGRAPH5\GRAPH5.EXE ^.gra
Q98=C:\APPS\QUICKTAX\QTAX98.EXE ^.Q98
Q97=C:\QCKTAX97\QTAX97.EXE ^.Q97
Q99=C:\QUICKTAX\Qtax99.exe ^.Q99
```

This sample illustrates common items that you will find in your `win.ini` file. `run` lines automatically load applications when the shell loads. `device` lines in the `win.ini` file make hardware devices available for older applications that need to use them. Wallpaper and desktop settings are here as well. Finally, for more backward compatibility, `wini.ini` contains a mapping of some of the registered file extensions.

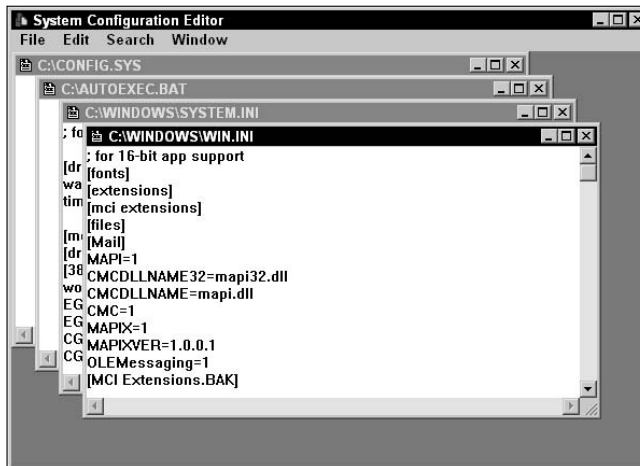


Corruption or deletion of the `win.ini` file tends to be less noticeable than that of the `system.ini` file, but it will still cause problems with older applications that refer to the information that is found in the file. If you have a problem with an older application, you should be able to fix the problem by re-installing the application.

sysedit.exe

The *System Configuration Editor* (shown in Figure 6-2), which you open by typing `sysedit.exe` at a command prompt, gives you quick access to your main configuration files — `win.ini`, `system.ini`, `protocol.ini`, `config.sys` — and to `autoexec.bat`. It is capable of editing and saving these files. This task could have been accomplished with any text editor, but it is far easier to just type `sysedit`.

Figure 6-2:
sysedit.exe
makes
getting at
configura-
tion files
easier.



Getting an A+

This chapter examines the process that is undertaken by the software on your computer during the boot process. The basic boot process for both Windows 9x and Windows XP computers is discussed, as well as the differences between them.

Key points that you should remember about this chapter are:

- ◆ Conventional memory used by MS-DOS applications is limited to 640KB, but can be expanded with `emmm386.exe` and extended with `himem.sys`.
- ◆ Virtual memory is managed by the Virtual Memory Manager, which presents a 4GB address space to applications on the system.
- ◆ Virtual memory is made up of physical RAM and hard drive space in the form of a swap file or paging file.
- ◆ The Windows 2000 and Windows XP boot process uses the `ntldr`, `ntbootdd.sys`, `ntdetect.com`, and `boot.ini` files.

Prep Test

- 1** A user has a program that automatically starts up when he logs into his computer. He would like to disable it. Where should he look for possible settings?
- A load = lines in win.ini
 - B The Startup group in the Start menu
 - C The RUN key in the Registry
 - D All of the above
 - E None of the above
- 2** A user has attempted to boot his computer, but instead of seeing the Windows XP boot menu, he sees a message stating, "NTLDR is missing. Press any key to restart." What should he do first?
- A Set the CD-ROM as bootable in the system CMOS and insert the Windows XP CD. This will allow him to complete an emergency repair.
 - B Reinstall Windows XP.
 - C Reboot the computer and enter Safe Mode, then choose RebuildBootSect.exe from the \windows\system32 directory.
 - D Remove the floppy disk from the A: drive.
- 3** The search order for bootable drives is stored in which location?
- A System BIOS
 - B PNP Configuration
 - C io.sys
 - D CMOS Memory
- 4** What driver is used to access virtual memory?
- A ramdrive.sys
 - B himem.sys
 - C emm386.exe
 - D None of the above
- 5** If you need to quickly edit win.ini, autoexec.bat, and config.sys, what command(s) could you use? (Select all that apply.)
- A sysedit.exe
 - B cfgedit.exe
 - C msconfig.exe
 - D notepad.exe win.ini autoexec.bat config.sys

Working with System Files and the Boot Process

- 6** What is the first file that is loaded as part of the Windows XP boot process?
- A io.sys
 - B ntbootdd.sys
 - C ntbooter
 - D ntldr
- 7** How large is the memory address space that is presented in applications running on a Windows XP computer?
- A 640KB
 - B 1MB
 - C 1GB
 - D 4GB
- 8** How large is conventional memory?
- A 64KB
 - B 640KB
 - C 1MB
 - D 1GB
- 9** What file is a copy of the SCSI drivers for the SCSI controller in your Windows XP computer?
- A ntbootdd.sys
 - B boot.ini
 - C ntldr
 - D ntdetect.com
- 10** What is name of the file that Windows XP uses to store memory data for the virtual memory system?
- A Memory file
 - B Swap file
 - C Storage file
 - D Paging file

Answers

- 1** D. Automatic commands can be found in any of the listed locations (`win.ini`, Startup group, or the Registry). *Review "Standard Boot Process for Windows XP"*
- 2** D. `ntldr` is missing in the error message that you will see if you attempt to boot from a Windows XP formatted disk that does not contain `ntldr`. *Check out "Standard Boot Process for Windows XP"*
- 3** D. CMOS memory contains the boot device order. *Take a look at "Power-On Self-Test (POST) process."*
- 4** D. Virtual memory reserves space on the hard drive to be used as RAM. It is managed by VMM32 and is loaded by Windows. *Peek at "Virtual memory."*
- 5** A, C. `sysedit.exe` and `msconfig.exe` provide quick means of editing the listed configuration files. *Refer to "Examining Other Boot Process Files."*
- 6** D. `ntldr` is the first file that loads as part of the Windows 2000 boot process. *Examine "Standard Boot Process for Windows XP"*
- 7** D. The memory address that is present to each application or VM is 4GB in size. *Examine "Virtual memory."*
- 8** B. Conventional memory is the first 640KB of RAM on a computer. *Review "Conventional memory."*
- 9** A. The `ntbootdd.sys` file is a copy of the SCSI driver that has been renamed and is used during the boot process. *Refer to "Standard Boot Process for Windows XP"*
- 10** D. Windows XP refers to the file used by virtual memory as a paging file or the page file. *Look over "Virtual memory."*

Book VI

Managing the Operating System

The 5th Wave

By Rich Tennant



“You know, this was a situation question on my A+ Certification exam, but I always thought it was just hypothetical.”

Contents at a Glance

Chapter 1: Loading and Configuring Device Drivers	681
Chapter 2: Working with Applications	699
Chapter 3: Optimizing the Windows Environment	735
Chapter 4: Using Windows-Based Troubleshooting Utilities	759

Chapter 1: Loading and Configuring Device Drivers

Exam Objectives

- ✓ Requirements for and the installation of device drivers onto a computer
- ✓ Identification and use of signed drivers
- ✓ Working with drivers and diagnosing driver errors

The United Nations' goal is to allow the nations of the world to work together and achieve common goals. They achieve this end by communicating with each other — but first they have to overcome the language barriers, which is done by employing hundreds of translators. *Device drivers* serve a similar role for computers that translators serve for the UN, converting information between the operating system and various pieces of hardware. Device drivers are responsible for establishing and maintaining communication links between the operating system and the various hardware devices that exist on the computer. Each hardware device has its own language. The job of the device driver is to translate data from the device into something that the OS can understand, and then translate the OS data into something that the device can understand.

Even with the advent of Plug and Play, the A+ Certified Professional still needs to understand how drivers work and how to diagnose driver-related problems. Even though Plug and Play may allow for default drivers to be automatically selected and loaded, the entire process of managing drivers is not a hands-off process, and intervention will be required on many occasions.

Loading and Removing Device Drivers

Because device drivers are so important, you should consider how device drivers are loaded onto your system. In the following sections, you take a look at loading device drivers for Windows 2000 and Windows XP. In Book III, you see how to configure the hardware to be added to a system; in this chapter, you focus on the software side of configuring these devices.



Because so many devices and drivers are floating around out there, the steps in this chapter for working with and managing drivers are necessarily generic.

Drivers and their installation are tied to `.inf` files, which contain installation instructions, and there are several ways to install drivers with `.inf` files:

- ◆ You can use Add/Remove Hardware in Windows 2000 or Add Hardware in Windows XP.
- ◆ You may be able to use the `.inf` files themselves.
- ◆ You may have to load the drivers manually.
- ◆ You may be able to install drivers directly through a setup program from the manufacturer, which may be used to install other software as well.

To remove drivers, you have the option of using Device Manager. The proper and clean removal of devices represents a large step forward in device management for the Windows operating system.



Before attempting device detection, ensure that the device is properly plugged in and powered; if it is not, the device will never be detected. If the device is not detected, then you will have to manually load the drivers and configure the device. Loading the drivers manually is discussed later in this section.



Some hardware vendors recommend that you copy their drivers onto your system before connecting the hardware. This process copies the drivers to your system and updates the `.inf` files to acknowledge that this hardware exists. This then allows the device to be properly detected when Windows finds it and allows Windows to automatically load the drivers. This information is usually included with the vendor's "Quick Install" or "Quick Start" guide. In most cases, it does not really matter when the drivers are copied — if the drivers are not on your system, then Windows will ask you where to find them.

To install a device driver, follow these steps:

- 1. Open the Add Hardware applet by choosing Start⇨Control Panel⇨Printers and Other Hardware⇨Add Hardware.**

The Add Hardware Wizard starts.

- 2. Click Next.**
- 3. Indicate whether you have already added the hardware to your system and click Next.**

If you are adding a device, Windows performs a search of Plug and Play devices. The results of this search are displayed for you and include any unknown devices.



Because adding a device also includes a troubleshooting option, this is your opportunity to select a device already installed on the system that you would like troubleshoot. If you select to troubleshoot a device, the next screen in the wizard informs you that when you finish the wizard, you will start the troubleshooter. The troubleshooter lets you set resource settings, such as IRQ and I/O addresses.

4. If the device that you need to load the drivers for is not listed, choose the option to Add a New Hardware Device at the bottom of the list and then click Next to proceed.

5. You have the option of letting Windows search for new hardware (shown in Figure 1-1), or you can select the hardware from a list.



If your computer appears to not be responding, give it some time to finish its detection process.

Figure 1-1: Plug and Play detection is used to initially locate all devices on your computer.



6. If no devices are found, click the Next button to manually choose a driver.

7. To install a device manually, first choose the hardware type (shown in Figure 1-2). After you make your selection, click Next.

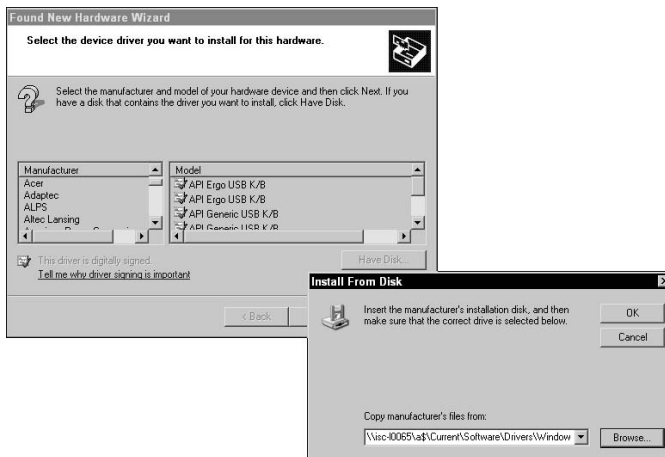
8. A screen listing manufacturers and models appears. Choose the manufacturer and then choose a model from that manufacturer. Click Next after making your selection.

If the device cannot be matched to known devices, you can click the Have Disk button to browse for or type the location of the device's .inf file, as shown in Figure 1-3.

Figure 1-2: Choose the type of device you are installing so that you can look at a filtered list.



Figure 1-3: Device lists are broken down by manufacturer to make locating devices easier.



If Windows is unable to detect the device settings for your device, it asks you to manually enter the settings for some or all of the hardware resources, which are I/O addresses, Memory, DMA, and IRQs.

You may be prompted for special device-related settings, but, if not, you should have a hardware installation dialog box on your screen.

9. Click Next to start installing your hardware.

If Windows cannot locate some of the required files, it may ask you to insert a disk with the drivers on it or to locate the files that it needs.

10. Click Finish when the installation is complete.

If you need to manually adjust the hardware resource settings, you can change the settings for the device by using either the device's control panel (if one exists) or the Device Manager.



Because control panels for each device are different, I look only at the Device Manager in this section.

To access the resource settings through Device Manager, use these steps:

1. Choose Start→Control Panel, open the System applet, click the Hardware tab, and then click the Device Manager button.

The Device Manager opens, and you can expand the device tree and select the device that you want to work with.

2. When you find the device, right-click it and choose Properties.

The Properties window for the device lets you troubleshoot the device, examine driver details, change resource settings, and choose which hardware profiles will use the device.

3. Click the Resources tab to change the resource settings that are used by the device, as shown in Figure 1-4. After adjusting the settings, click OK to close the device properties dialog.

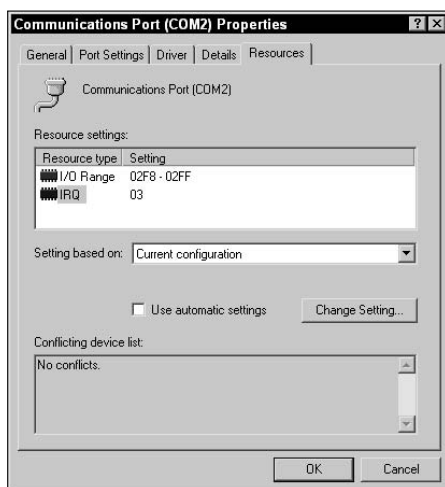


Figure 1-4: You may have to tell Windows XP what resources your device requires.

At some point, you may want to remove a device. To do so, you can use either the Add/Remove Hardware Wizard in Windows 2000 or the Device Manager with all versions of Windows. The Device Manager allows you to select any installed device on your system and delete it.



Device Manager is the main tool used for configuring and removing installed device drivers.

To use the Device Manager to remove a device, follow these steps:

- 1. Open the Device Manager by choosing Start→Control Panel, open the System applet, click the Hardware tab, and then click the Device Manager button.**
- 2. Expand the hardware tree to locate your device, right-click the device, and choose Uninstall.**

You are presented with a confirmation dialog box.

- 3. Click OK and then the device will be uninstalled.**

If you use the Add/Remove Hardware Wizard in Windows 2000, you can uninstall a device or unplug/eject a device, as shown in Figure 1-5.

Figure 1-5:
Device removal now comes with options. You can temporarily remove devices with the unplug option.



Unplug/Eject a device can be used to temporarily remove a device (such as a PC card) that you know you will be reinstalling later. If you are planning to permanently remove the device, choose the Uninstall a Device option and then click the Next button. You are presented with a list of devices; select the device that you want to uninstall and click Next. Confirm that you want to delete the device by selecting Yes and then clicking Next. The removal of the device is completed. The next screen tells you that the device has been uninstalled; click the Finish button.

Changing Device Settings

Regardless of how you install the drivers for the device, if Windows can't detect the settings for the device, you are prompted to specify IRQ, I/O, DMA, and memory resource settings. The prompt should only be necessary for non-Plug and Play devices. If you are not sure of the current settings for the device, confirm the settings by examining the hardware or by using the configuration utility that came with the device. A configuration utility should be included on the driver disk of any devices that support software configuration. If the device is configured for Plug and Play, you shouldn't have to worry about this.

For many devices, you can add drivers from the device's Control Panel applet, if there is one. This ability is not available for all devices but is available for modems and several other devices.

If you want to change the settings for a device that is installed on your system, you can use either the Control Panel applet for the device (such as Phone and Modem Options, Display, or Network) or the Device Manager. Because each Control Panel applet is different, I don't go through each of them here, but rather show you the Device Manager.

To access the Device Manager, choose Start→Control Panel, double-click the System applet (opening the System Properties dialog box), and then click the Device Manager button on the Hardware tab. You can also open the System Properties dialog box by right-clicking My Computer and choosing Properties. With the Device Manager open, expand the appropriate device tree until you see the device that you want to work with; then right-click the device and choose Properties. You can use the Properties window to disable the device in the current hardware profile or all hardware profiles, change the driver, and set resources for the device (see Figure 1-6). Based on how the driver was designed, you might not have a Driver or Resources tab and won't be able to configure the settings on those tabs.

If you want to delete a device driver that has been installed, you can do this in the individual device's Control Panel applet, such as the Network Control Panel applet, or you can use the Device Manager. If you use the Device Manager, you can delete any device that is listed in the computer. With the Device Manager open, expand the device category your driver belongs to, select the device that you want to delete, and then click the Delete button. As with the Uninstall/Unplug options in Windows 2000, Windows will, by default, unplug devices that you have removed, leaving drives and settings behind for you in the event that you want to reinstall the device in the future. To see devices with Device Manager that have only been unplugged, choose View→Show Hidden Devices in the Device Manager. This now shows all devices that have been installed on your computer, even if they are currently unplugged, and you can choose to remove them here and delete their drivers.

Figure 1-6:
Device
Manager
is a very
powerful
tool when it
comes to
device man-
agement.



Lab 1-1 will guide you through examining the resources that are in use by your network card. In this exercise, you can use either Windows 2000 or Windows XP. Lab 1-1 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Updating Drivers

From time to time, software vendors update their drivers, and you can choose to not update them, manually update them, or use Windows Update. When these drivers make it to the Windows Update Web site or Microsoft Update Web site, if you have upgraded, you can add them to your updates by choosing the optional hardware components. If you choose to update the drivers manually, you right-click the device in Device Manager and choose Update Driver or click the Update Driver button on the Drive tab. Because some drivers make it to the Windows Update Web site, the first question you will be asked when you choose to update your driver is whether you want to check Windows Update for copies of the drivers. If you choose not to use Windows Update or if there are no updated drivers on the site, you can still update them manually.

The next step in the update process is either allowing Windows to automatically locate the drivers (Recommended) or having you manually specify a location for the drivers (Advanced). If you use the automatic method, Windows checks your currently installed `.inf` files and the root directories of your removable media. If no matching drivers are found, then you need to come back and choose Advanced, which allows you to recheck the same

paths as well as specify a directory in which to look for drivers. If all else fails, you can choose “Don’t search. I will choose the drives I want to install.” This option gives you a list of all drivers as well as a Have Disk button to load other drivers into the system. In most cases, you won’t have to use this screen. Just telling Windows to search a specific directory is usually enough for it to locate the correct driver for you.

Signed and Unsigned Drivers

Driver signing was incorporated into Windows 2000 to prevent the installation of untested drivers on the system. Hardware vendors have an option to send their completed drivers to Microsoft for testing. After testing the driver and confirming that it doesn’t have any apparent problems within the OS, Microsoft returns the driver to the vendor with a signature file. This signature file has a .cat extension and contains information about the original driver file to ensure that the driver was not modified since it was sent to Microsoft. When loading drivers, Windows 2000 and Windows XP check for the existence of the signature file. If the signature is not available, then you can configure what you want to have happen on the system. Your choices are:

- ◆ Ignore the missing signature and install the drivers.
- ◆ Warn the user of the missing signature, but still give the option of installing (see Figure 1-7).
- ◆ Block the installation of the driver.



Figure 1-7:
Driver Signing prevents the installation of unwanted drivers.



These settings are configured in the System control panel, using the Driver Signing button on the Hardware tab. If you are logged in as an administrator, you can save these settings to affect all users. Typically, the option to block unsigned drivers is enforced on high-availability workstations or servers.

Any individual user has the ability to make his or her individual driver signing policy stronger than the administrator's defined system default.



Driver signing is designed to increase system stability, by requiring hardware manufactures to validate their drivers with Microsoft.

Working with Plug and Play

Thanks to Plug and Play, you have to do less work to add a device because device identification, driver selection, and hardware configuration are handled by Plug and Play (at least once the device has been added to the driver database). The Plug and Play specification says that a device should do the following:

- ◆ Be uniquely identified.
- ◆ State the services that it provides and the resources that it requires.
- ◆ Identify the driver that supports it.
- ◆ Allow software to configure it.

The Plug and Play standard also requires that Plug and Play devices be backward-compatible with non-Plug and Play systems, or at least that's the recommendation. All Plug and Play hardware should also be independent of operating systems that are installed on the system.

The Plug and Play process starts with the computer BIOS that supports Plug and Play. In addition, you also need an operating system that supports Plug and Play and Plug and Play devices. Windows 2000 and Windows XP are both Plug and Play-aware operating systems; Windows NT 4.0 does not support Plug and Play. The Plug and Play process contains three major components:

- ◆ Bus Enumerator
- ◆ Configuration Manager
- ◆ Resource Arbitrator

Figure 1-8 shows how these components fit together, and I discuss the whole Plug and Play process in the following paragraphs.

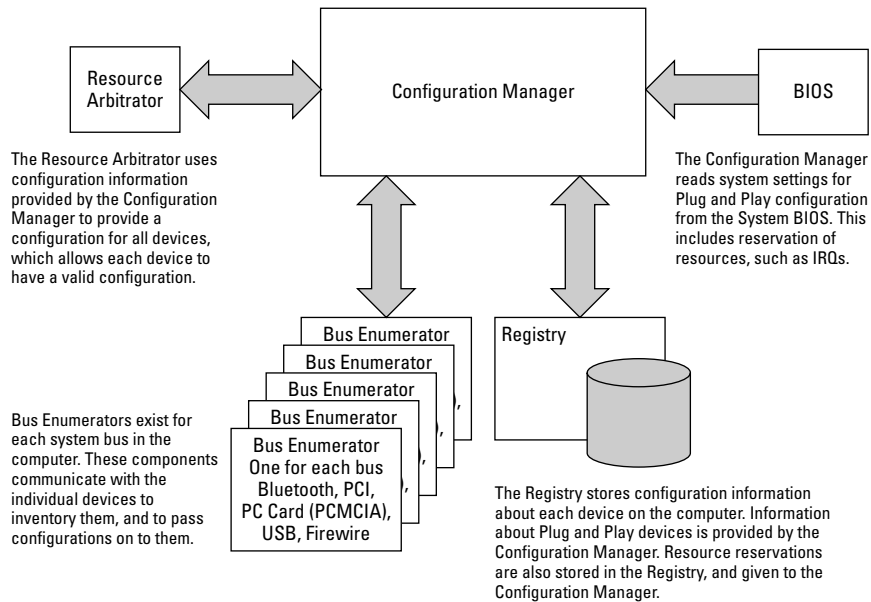


Figure 1-8: The Plug and Play concept is very simple in design.



When Windows NT 4.0 was released, it provided a shell (or user interface) update to make it appear like Windows 9x. There were also some architectural changes to provide additional power and stability over Windows NT 3.51. Because Plug and Play represented a major structural change in the OS, Microsoft chose not to support it in Windows NT 4.0.

At the lowest level in the Plug and Play process is the *Bus Enumerator*. The Bus Enumerator's job is to deal directly with individual devices. Each bus architecture in your computer has a Bus Enumerator, and the major bus architectures include: PCI, ISA, EISA, PCMCIA (PC Card), SCSI, and USB. When a device is added to the system, especially in the case of PCMCIA devices, it announces its presence to the Bus Enumerator. The Bus Enumerator assigns an ID to the device, interrogates the device IRQ and I/O preferences, and possibly reconfigures the device based on requests from the Configuration Manager.

The Bus Enumerator is called into action when a Plug and Play computer boots, and it has work to do whenever a device is added or removed from a particular bus. When the Bus Enumerator (and the rest of the Plug and Play system) is given a notification of device removal, it can prepare for it. For example, if you want to remove a PCMCIA Ethernet adapter, you can notify the Plug and Play system by stopping the device by clicking on the PC Card icon in the System Tray and choosing the Ethernet adapter.



Devices that can be removed or added to a system while it is running are referred to as *hot swappable*. *Cold swapping* takes place when the machine is powered down, which is the case for most devices. PCMCIA or PC Cards support hot swapping, and some of the newer computers support hot swapping on the PCI bus. Some manufacturers also use the term *warm swappable* to refer to components that can be replaced when the computer is in a sleep, standby, or other lower-power state.

The Bus Enumerator does *not* decide what resources the new device is going to use — that is the job of the *Configuration Manager*. The Bus Enumerator passes all of the information about the device to the Configuration Manager, which takes a look at all of the resources that are required by all of the devices on the system. If multiple devices require the same resources, the Configuration Manager calls on the Resource Arbitrator. The *Resource Arbitrator* evaluates which devices can operate at which settings and provides a solution in which all devices are able to use compatible settings.

After the Resource Arbitrator has worked its magic, a new resource configuration is passed to the Configuration Manager. The Configuration Manager accepts the new configuration; it contacts each Bus Enumerator and has the Bus Enumerator reconfigure any devices that require reconfiguration. So, although the Bus Enumerator doesn't decide how to reconfigure devices, it does actually carry out the reconfiguration of the devices.

All of the information about devices on your computer is written into the *Registry*, which maintains the hardware tree. The Configuration Manager is responsible for writing the information into the Registry. In a computer that supports the Plug and Play BIOS, the basic Plug and Play configuration is actually performed at the BIOS level. When the operating system loads, it goes through the Plug and Play process and modifies any configurations that require modification.

If legacy devices (non-Plug and Play devices) exist on your system, resources are assigned to them first. This makes perfect sense because these devices cannot be reconfigured through the Plug and Play process. In order to ensure that the proper resources are allocated to legacy devices, you can reserve those resources in the system BIOS.

In most BIOS configurations, you have the option of reserving resources for non-Plug and Play devices. The settings in the BIOS tell you whether you are reserving resources for non-Plug and Play devices, reserving resources for legacy devices, or disabling resources. Those resources are not actually disabled, but their use by Plug and Play devices is disabled; legacy devices are still able to use these resources. The BIOS represents the best place to reserve legacy resources for this reason: It represents the first area where Plug and Play resources are actually assigned to devices.



Before changing any of your BIOS settings, consult your motherboard documentation. Incorrect modification of your BIOS can result in your computer not functioning.

Because legacy devices cannot be reconfigured through the Plug and Play process, they should always be allowed to use the resources that they require. In a perfect world, that would be all you need. But we don't live in a perfect world, so you must reserve resources. Reserving the resources for legacy devices at the BIOS level is more consistent for achieving a working configuration than allowing them to be reserved at the OS level.

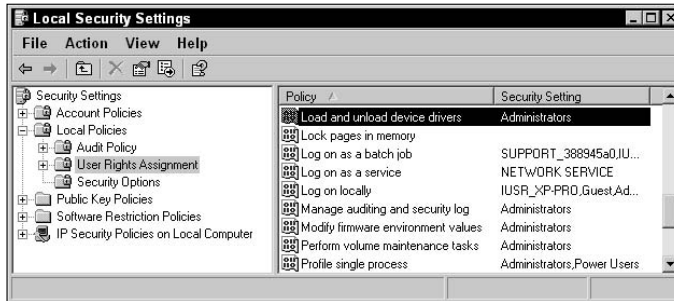
When Plug and Play systems first entered the market, they were plagued with difficulty. Much of this difficulty came from devices that were not fully Plug and Play-compliant or from non-Plug and Play-aware system BIOS, not to mention mixing Plug and Play and non-Plug and Play devices together in a computer, which always causes additional problems and was often the case in early systems. Over the years, hardware and software manufacturers have learned a lot about how Plug and Play should and can work, and today Plug and Play is much closer to its ideal.

Permissions to Install Drivers

In order to install drivers on a Windows computer, by default you need to be a member of the Administrators group. In most cases, you don't want just anybody adding hardware devices to a computer and causing no end of problems for the people responsible for maintaining the computer. This ability to install the software for the device is controlled in the same way all rights on the computer are controlled, through the Security Policy. The computer policy can be changed locally, but, if the computer is part of Active Directory, then conflicting Active Directory settings would override the local Security Policy. To view the active security settings on your computer and to see if they are being overridden, you only have to look as far as the Administrative Tools folder, located in the Control Panel on your computer.

In Administrative Tools, you will find the Local Security Policy application, which lets you see all of the security settings on your computer. To see the right to load and unload device drivers, expand the tree on the left to see Local Policies⇒User Rights Assignments. After selecting User Rights Assignments, you can browse the list on the right for Load and Unload Device Drivers. The Security Setting column tells you who has the permission, as shown in Figure 1-9. If you want to change the permission, double-click the security setting and the Load and Unload Device Drivers Properties dialog box will allow you to change the list of people who have permissions.

Figure 1-9:
The Local Security Policy allows you to install device drivers.



Verifying Driver Installation

To ensure that the drivers for a device are working properly, you should check three things:

- ◆ **The device is listed in Device Manager, and it displays a “working properly” message.** If you see the Question mark category, then you have some devices installed that do not drivers.
- ◆ **Is the device available to its own management or control software?** For instance, does the scanner show up in the scanning software application? If not, then it could be due to a faulty driver.
- ◆ **Is the device yielding predictable results when used?** If not, the issue could be driver-related.

In most cases, if the drive is listed as working in Device Manager, then the device driver is loaded properly for the device; the issues may actually be settings for the device or a hardware failure.

Device Driver Failure

What do you do if a device driver does not work? This is a question that’s heard many times by users with driver problems, and the answer is, as always with computers, “it depends.” Some actions that you can take to deal with a malfunctioning driver include:

- ◆ Upgrading to a newer driver version
- ◆ Removing the device
- ◆ Rolling back drivers
- ◆ Booting to the last known good configuration
- ◆ Using restore points

If the driver is not working or if it has compromised the stability of the computer, you do have the preceding options. Issues with drivers could range from the driver just not working to the computer having a Stop event (blue-screen error) during the boot process. In most cases, if the driver is just not working, then you look at how the device is connected to the computer and also examine any related components. So, if it is a USB device, are the rest of the USB components, like USB hubs and controllers, on the computer working?

When you have verified this, and that the device is properly attached and powered, then you should examine the driver version. Check for a newer (usually) driver or an older (rarely) driver that could fix your problem. Discussion boards often help diagnose driver-related issues. In most cases, you should initially load the component with the latest driver available, which means ignoring the CD that came with the device and going directly to the manufacturer's Web site. If the device doesn't work because of a hardware conflict, then you can simply remove the device and replace it with a different manufacturer's version.

Rather than just not working, the problem with the driver failure may be more prominent, such as a blue-screen error during system boot. If that is the issue, then your resolution is somewhat more critical. The first step is deciding what your immediate resolution needs to be. If it is to allow the computer to be usable and that the device was considered an extra feature, then remove the device and look at adding and troubleshooting the device installation at a later time. If it is to get the device working, then you will need to troubleshoot the device installation immediately. Your steps may vary some, but usually you want to do the following:

- ◆ **Boot into Safe Mode to verify that base system files are still operational.** From there, you can remove or disable the device to allow the computer to boot. This can be done while leaving the device in the computer.
- ◆ **If the error occurred as part of a driver update, choose the driver rollback option on the Drive tab of the device's Device Manager properties.**
- ◆ **Download a newer version of the driver; you might be trying to load an older version of the driver that has known problems.** The vendor's Web site will let you know if there are newer versions of both drivers and firmware and often will tell you what issues they corrected.
- ◆ **Risking the loss of configuration settings or files, the last two options are last-chance efforts to restore system stability in the form of booting to the Last Known Good Configuration or using a past system restore point.** Because these two items return to a past point in time, you do run the risk of losing additional settings that you were not aware of. For more information about these two procedures, see Book VII, Chapter 3.



In a perfect world, one of these fixes will work for you, and you won't lose any data or settings.

Getting an A+

In this chapter, you look at general processes to load and configure device drivers. Specifically, you find out the following:

- ◆ The driver installation procedures are found in `.inf` files.
- ◆ Drivers verified and signed by Microsoft include a `.cat` file, which contains the signature.
- ◆ Plug and Play takes care of device resource settings for most devices.
- ◆ Device Manager is the basic management tool for driver management.
- ◆ Driver errors can be recovered by updating drivers, using Safe Mode, loading the Last Known Good Configuration, or restoring to a system restore point.

Prep Test

- 1 Device drivers are used to do what?**
 - A Provide power to the device while it is operating.
 - B Act as a software interface between the operating system and the device.
 - C Provide advanced configuration options for the device.
 - D Steer the device to a new location on the computer.

- 2 What is the risk of using an unsigned driver?**
 - A Unsigned drivers are only reliable when using Windows NT 4.0.
 - B Unsigned drivers are missing their .sig files.
 - C Unsigned drivers are not backward-compatible with older hardware versions.
 - D Unsigned drivers have not been tested by Microsoft.

- 3 Which of the following is *not* part of the Plug and Play procedure?**
 - A Resource Arbitrator
 - B Dynamic Reconfigurator
 - C Configuration Manager
 - D Bus Enumerator

- 4 Which users on a Windows XP computer have the right to load device drivers?**
 - A Guests
 - B All Users
 - C Administrators
 - D Power Users

- 5 Device settings, such as IRQs, can be modified through which tools? (Choose all that apply.)**
 - A Device Manager
 - B Resource Manager
 - C Device's Control Panel, such as Network Control Panel
 - D Device Arbitrator

- 6 Which tool can be used to examine the status of a device?**
 - A Device Manager
 - B Resource Explorer
 - C Device Explorer
 - D Status Examiner

Answers

- 1 B. Device drivers act as an intermediary between the operating system and physical devices. *See "Loading and Removing Device Drivers."*
- 2 D. Unsigned drivers have not been tested by Microsoft's hardware qualification labs and, as such, are missing their .cat files. These drivers may work, but are not verified by Microsoft. *Review "Signed and Unsigned Drivers."*
- 3 B. The Bus Enumerator identifies and configures devices, the Configuration Manager coordinates devices, and the Resource Arbitrator assigns resources such as IRQs and I/O addresses. *Check out "Working with Plug and Play."*
- 4 C. Only members of the Administrators group have been granted the right to load and unload device drivers, by default. *Peruse "Permissions to Install Drivers."*
- 5 A, C. Most device settings can be modified through either the Device Manager or the device's control panel. Resource Manager and Device Arbitrator are made-up terms. *Take a look at "Changing Device Settings."*
- 6 A. Device Manager can be used to examine the status of the device. *Peek at "Verifying Driver Installation."*

Chapter 2: Working with Applications

Exam Objectives

- ✓ **Identify how MS-DOS, 16-bit Windows, and 32-bit Windows applications operate in Windows 2000 and Windows XP**
- ✓ **Diagnose applications install, start, or load failures with Windows XP**

When MS-DOS was first created, there were very few applications for it. Many people chose to use the BASIC programming language to write their own applications. But in short order, people developed faster applications by using first Assembler and then higher-level languages like C. These applications served users well for several years, and when Windows first arrived on the scene, Microsoft built in as much backward compatibility for these older applications as it could. Microsoft knew that if the Windows operating system (OS) did not allow older applications to work with it, current MS-DOS users might not want to adopt it. To this day, Microsoft attempts to provide the most backward compatibility possible in its operating systems.

When the 16-bit Windows environment, such as Windows 3.0, needed applications, they were developed and worked well. In later generations, these applications were superseded by 32-bit versions for the even newer Windows 9x operating system. The newer applications worked with data in larger chunks and ran faster. In this chapter, you take a look at how all of these applications run on your computer.

In order to properly maintain and diagnose applications, an A+ Certified Professional must be able to manage the installation and removal of applications. In addition to these skills, this chapter also reveals the operating architecture in which applications run. Understanding this architecture allows the professional to diagnose problems quickly. You also examine the level of backward compatibility that is built into the modern Windows operating systems, and you explore multitasking.

Installing and Removing Applications

Before you are able to work with applications on your computer, you need to install them. With the rate at which the computer industry changes, you cannot avoid the need to upgrade or remove applications on your computer as they become obsolete.

Working with .msi Files

With Windows 2000, Microsoft created a new setup installation file. .msi files are associated with `msiexec.exe` or the Windows Installer service. `msiexec.exe` is a generic application installer, and the Windows Installer service can be used to push applications out to computers on your network through Active Directory.

The .msi file contains the installation script, and in most cases also has a compressed copy of the files that need to be installed. If they are not stored in the .msi file, then you will also have a .cab file.

Windows Installer or .msi installations have these benefits:

- ✔ Consistent user interface and command-line interface

- ✔ Standard and easy way of customizing applications
- ✔ Diagnose and repair configuration problems at application run-time
- ✔ On-demand installation
- ✔ System-wide management of shared resources
- ✔ Restore the pre-installation state of a computer in the event of installation failure
- ✔ Elevated installation privileges

To install an .msi file, simply double-click on it and follow any installation instructions.



The terms *application* and *program* are synonymous; they are both programming code that performs a function. But in this chapter, I use the two terms to mean slightly different things. I use the term *application* to describe the programming code that represents the functions you want installed on your computer — like a word processing application or a game. I use the term *program* to describe the programming code that allows you to install the application.

Installing an application

Most applications come with an installation program that must be run in order to install the application. Some applications, like Kaufman Mail Warrior, are actually standalone applications that do not require extra files and settings to be created on your computer, but these are rare. Most applications require several files to be installed and often require Registry entries to be created to hold the settings for the application. Due to the complexity involved in copying the files and creating the settings, you use the installation program to ensure that the application is installed properly. The application's developer decides the name of the setup program, which in many cases is `setup.exe` or `install.exe`.

Very few installation programs work in exactly the same way. In general, though, you are asked for the location where you want the application installed. Other options, such as whether you want to create a desktop shortcut to the application or whether you want a specific option to be enabled, are application-specific.

Pre-Windows 2000

Starting with Windows 95 and, shortly thereafter in the Windows NT family with Windows NT 4.0, Microsoft began providing a Control Panel applet called Add/Remove Programs to manage the installation and removal of applications on your computer. To install a new program on a pre-XP computer, do the following:

1. Choose Start→Control Panel.
2. Double-click Add/Remove Programs.

The Add/Remove Programs Properties dialog box appears, as shown in Figure 2-1.



Figure 2-1:
The Add or Remove Programs Control Panel applet from Windows ME.

3. Click the Install/Uninstall tab if it isn't already selected.
4. Click the big, friendly Install button.

This opens the Install Program from Floppy Disk or CD-ROM wizard, which prompts you to insert the floppy disk or CD-ROM that has your installation program.

5. Click the Next button.

The Control Panel applet scans your floppy disk and CD-ROM drives for any programs named `install.exe` or `setup.exe`. If either of these programs is found, it is launched. Otherwise, you're asked to provide the command path for the installation application.

6. If necessary, provide the path to the installation application.

7. To start the installation, click, ironically, the Finish button.

Windows 2000 and later

The format of the Control Panel applet changed with Windows 2000, while the name remained as Add/Remove Programs. When Windows XP was released, the name was changed to Add or Remove Programs. Installing a new program on a Windows XP computer is slightly different from Pre-Windows 2000 versions of the OS. Your course of action looks like this:

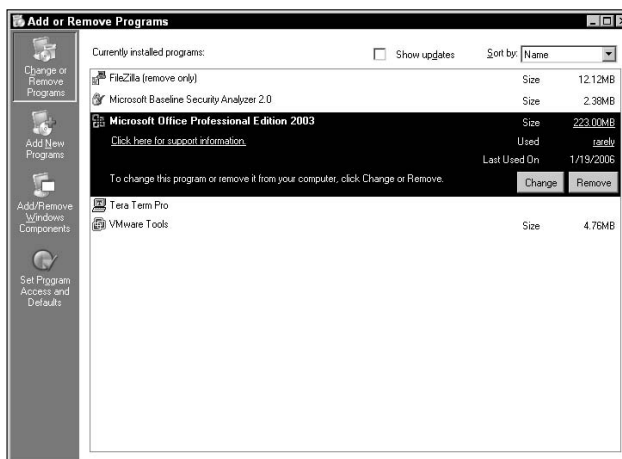
1. Click Start→Control Panel.

2. Double-click Add or Remove Programs.

Depending on how your interface is set up, Add or Remove Programs might appear in a drop-down list attached to the Control Panel, in which case you just single-click Add or Remove Programs in the drop-down list.

The Add or Remove window appears, as shown in Figure 2-2. Notice the toolbar along the left edge of the window.

Figure 2-2:
The Add or Remove Programs Control Panel applet in Windows XP. This format began with Windows 2000 and is also used with Windows Server 2003.



3. Click the Add New Programs button on the left-side toolbar.

The window adjusts itself to your installation needs.

4. Click the CD or Floppy button at the top of the window.

This opens the same Install Program From Floppy Disk or CD-ROM wizard that you saw in the step list for Pre-Windows 2000 computers.

5. Click the Next button.

Windows XP scans your floppy disk and CD-ROM drives for a program named `install.exe` or `setup.exe`. If either program is present, it is executed. Otherwise, you are asked to provide the path to the setup program for the application you want to install.

6. If necessary, provide the path to the setup program and then click Finish.

If your computer is part of an Active Directory domain, then you may also see applications listed in the section of the Add or Remove Programs Control Panel applet called Add programs from your network. You can filter the list using the Category drop-down list, and install these by selecting an application and choosing the Add button. This will launch the installation program.

Removing an application

Many applications today are distributed on CD-ROM. When you insert a CD-ROM into your CD-ROM drive, Windows searches for an `autorun.inf` file and launches the program that is specified by the `autorun.inf` file. In most cases, the specified program is the application setup program if the application is not already installed. This “AutoPlay” feature of Windows makes installing applications even easier.

Uninstall and the Registry

If an application uninstalls incorrectly or you would like to know what program is used to uninstall an application, the settings are stored in the Registry under the `UnInstall` key, which is found at:

```
HKLM\Software\Microsoft\Windows\  
CurrentVersion\Uninstall\  
<Application_Key>\UninstallString
```

If the application was uninstalled but still shows up in the Add/Remove Programs Control Panel applet, then you can delete the entire `<Application_Key>` from the Registry.

With Windows XP, when you attempt to uninstall an application that is missing the uninstaller, it will ask if you want to remove the entry from your Add or Remove Programs Control Panel applet.

Most applications also provide a program to uninstall the application when you no longer want it on your hard drive. The path to the program and any switches that are required for it to function are usually stored in the Registry during the installation of the application.



Always use care when working in the Registry; improper changes can leave you in a position where you will be required to re-install the operating system.

The procedure to uninstall or remove a program from a Windows XP computer is very easy. After opening the Add or Remove Programs Control Panel applet, select an installed application from the list. You are then provided more details about that application, like its size and how often you use it. You will also see one or more buttons beside the application: Change, Remove, or Change/Remove. To change the installed components of an application, click the Change button, and to remove the application, click the Remove button. Some applications will have a Change/Remove button, which will then allow you to make changes to the list of installed components, or to remove the application; it uses the same program to accomplish both tasks. Windows looks up the name of the uninstallation program to run in the Registry and executes it.

Like the installation procedure, there is no set uninstall procedure, and it is left up to the software developer to design the uninstall routine. Some developers do an excellent job, while others do not. In the case of applications that do not uninstall properly, you may find the icons, files, or Registry entries are still present after the uninstall program has completed. If that is the case, you will have to manually remove the leftover components.



Make sure that you are familiar with the standard method of adding and removing applications from your system through the Add or Remove Control Panel.



In many cases, applications are removed by using special options with the setup program that originally installed them. One nice feature of using the setup program is that many developers allow you to not only remove the application, but also to change the components that are installed.

Getting the Most out of Multitasking

Multitasking is a computer's ability to balance the processing time given to multiple applications that are running at the same time. It comes in handy when running multiple applications at the same time.

There are two basic types of multitasking: *Cooperative* and *preemptive*. In this section, I discuss differences between the two types of multitasking, and by the end, you should have a clear picture of which one is better.

Cooperative multitasking

Cooperative multitasking means sharing the time on the processor by cooperation. By its name, you would think cooperative multitasking is the better form of multitasking. This is probably because cooperation has always been thought of as a good thing. If you needed to get a job done (like building a house) and you had five people who were able to help you, through teamwork and cooperation, you could get the house raised quickly. This is a good illustration of when cooperation works. However, if a couple of those people do not work with the rest of the team (they go to get more nails and never come back), their lack of cooperation could slow the entire process down.

With cooperative multitasking in the computer environment, a few programs that do not cooperate well with others can slow the entire process down. If you launch Microsoft Excel, for example, and start a large recalculation of the entire spreadsheet, Excel occupies 100 percent of the processor's time. At periodic intervals, Excel checks to see whether any other programs require processing time, at which time Excel turns control of the processor over to the other applications. Each of these applications follows the same process that Excel does; they occupy 100 percent of the processor's time if needed, and give up control only when they reach their periodic interval for checking with the rest of the operating system. If a program doesn't check with the other applications often enough, it is thought of as a non-cooperative program. The big problem with non-cooperative programs is that they can hog 100 percent of the processor time for as long as they want.

Cooperative multitasking was implemented to allow multitasking in a 16-bit Windows environment. Most of that environment was *single threaded*, which means that there is only a single element (or thread) in the operating system that gives the processor instructions. As such, all applications had to work while sharing a single thread of execution. It may seem that cooperative multitasking is inefficient, but cooperative multitasking is better than no multitasking.

Preemptive multitasking

Preemptive multitasking takes a different approach to multitasking in the Windows environment. With *preemptive multitasking*, the operating system decides which applications get execution time. Preemptive multitasking is designed to work in a 32-bit *multithreaded* environment, in which applications have multiple execution threads, and the operating system is capable of managing multiple threads issuing instructions to the processor.

Each application is given a certain percentage of execution time to use during each second. The operating system then manages each of the processes that access the processor. With the operating system acting as the conductor,

sharing of the processor is more equal with fewer conflicts. This does not mean that certain tasks do not attempt to run away with all of the processor's computing time; it is just less likely to happen. Preemptive multitasking entered the Windows operating system with the development of Windows 95 and Windows NT and has been improving ever since.

Running 32-Bit Windows Applications

The 32-bit Windows applications are the ideal type of applications to run under a Windows 32-bit environment. Windows 9x, Windows ME, Windows NT, Windows 2000, Windows XP, Windows Vista, and Windows Server 2003 all represent Windows 32-bit environments. Since all of these operating systems are 32-bit in nature, it only makes sense that the applications that are run in these environments are also 32-bit. With that said, you should understand the benefits of running 32-bit applications and how those applications are executed in Windows 9x and Windows XP environments.

Benefiting from 32-bit applications

There are several benefits to running Windows 32-bit programs, such as

- ◆ **Multithreading**
- ◆ **32-bit data transfers**
- ◆ **Process isolation**

One of the greatest benefits of 32-bit Windows applications is the ability to be multithreaded. *Multithreaded* applications run several threads of code concurrently. Each one of these threads usually is assigned to a specific task. In the case of Microsoft Word, different threads can process typed characters, check spelling, or check grammar in your document all at the same time. If your computer has only one processor, then only one task is actually performed at any given instance, but the scheduling of all of these different tasks is optimized. If you are lucky enough to have a multiprocessor computer, then each of these threads can be assigned to execute on a different processor. This not only optimizes the execution of the program, but also evenly utilizes the processors.

The 32-bit Windows applications work with data in 32-bit blocks, just as 16-bit applications work with data in 16-bit blocks. In any given clock cycle, a 32-bit application should be able to process more information than a similarly written 16-bit Windows application. In a perfect world, the speed factor of a 32-bit application would be twice that of a 16-bit application, but we do

not live in a perfect world, so this multiplier is not actually realized. Other factors that affect the performance of the application include how the logic in the code was optimized to allow these 32-bit blocks to provide better performance.



Microsoft now offers a 64-bit version of Windows XP Professional and Windows Server 2003 to run on computers with new 64-bit processors. Eventually, 64-bit applications will be written to take advantage of this faster architecture.

The 32-bit Windows applications also run with some level of isolation from other applications that are running on the system. This provides better stability for the applications that are executing. There are some differences between how this is actually accomplished in Windows 9x and Windows XP computers. These differences are examined in the next two sections.



Ensure that you can recall the benefits of using 32-bit Windows applications on your system, rather than 16-bit Windows applications.

Executing in the Windows 9x environment

The Windows 9x architectural diagram is shown in Figure 2-3. Note that each 32-bit Windows application runs in its own area, but still runs under the system virtual machine (VM). Each VM in Figure 2-3 emulates the physical hardware that is found in the computer, allowing the hardware resources to be shared between the VMs. Each of the 32-bit Windows applications has its own memory space, message queue, and potentially multiple threads. Having its own memory space provides a couple of basic functions:

- ◆ It reduces the risk of conflict.
- ◆ It is easier to terminate all resources in the event the application hangs or crashes.

The application memory space contains the executable application code as well as any DLLs (Dynamic Link Libraries) or other code that the application may have loaded into memory.

When running 32-bit Windows applications under Windows 9x-based OSes, you are provided with a very stable environment for executing your code. The only problem with this environment is that it still relies on an OS that shares some of its code with the 16-bit computing environment. This hampers the Windows 9x-based OS with some inherent system instability. Although system instability is the result, the OS was set up this way to increase backward compatibility.

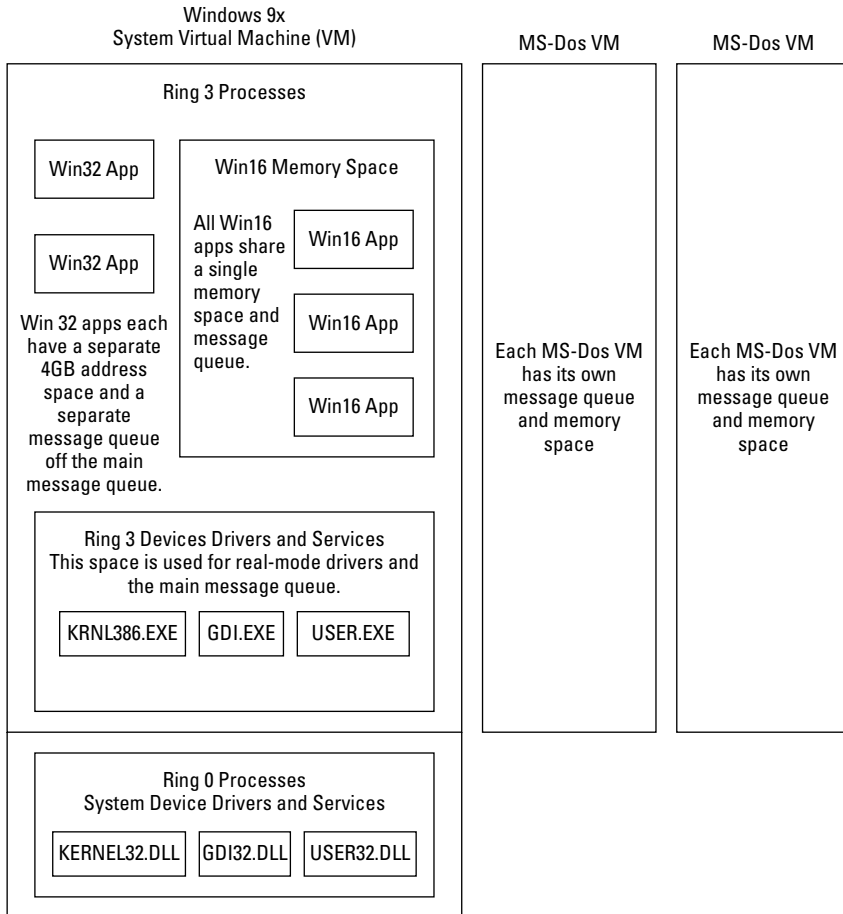


Figure 2-3: 32-bit applications run in a virtual machine that is shared with 16-bit applications.

Executing in the Windows 2000 and Windows XP environments

The architectural diagram of Windows 2000 is shown in Figure 2-4. Note that 32-bit Windows applications are executed as separate processes in the user mode portion of the operating system. This architecture is identical to that of Windows XP.

In Windows XP and all Windows NT-based OSes, 32-bit Windows applications are executed in an area that is completely separate from where the operating system executes. Unlike the Windows 9x operating system, the Windows XP OS does not actually contain any 16-bit code. This provides the

best stability for both the operating system and the applications that are running in it, but does hinder some of the backward compatibility of the product. This does not mean that Windows XP is not backward compatible, but rather that it is less backward compatible than Windows 9x. There have been attempts to increase the backward compatibility of the Windows NT-based product line, so applications that may not have worked with Windows NT 4.0 may work under Windows 2000, and applications that may not have worked under Windows 2000 may work under Windows XP.

Each 32-bit Windows application is given

- ◆ **Its own address space to work with:** If an application hangs or crashes, the separate memory address space makes it easy to destroy or flush the entire memory that deals with that application.
- ◆ **A separate message queue:** By having a separate message queue, if an application hangs or crashes, that halted application doesn't block the messages that are destined for other applications.
- ◆ **The ability to own several threads that are preemptively multitasked by the operating system:** Multiple threads per application allow the application to perform multiple operations at the same time.

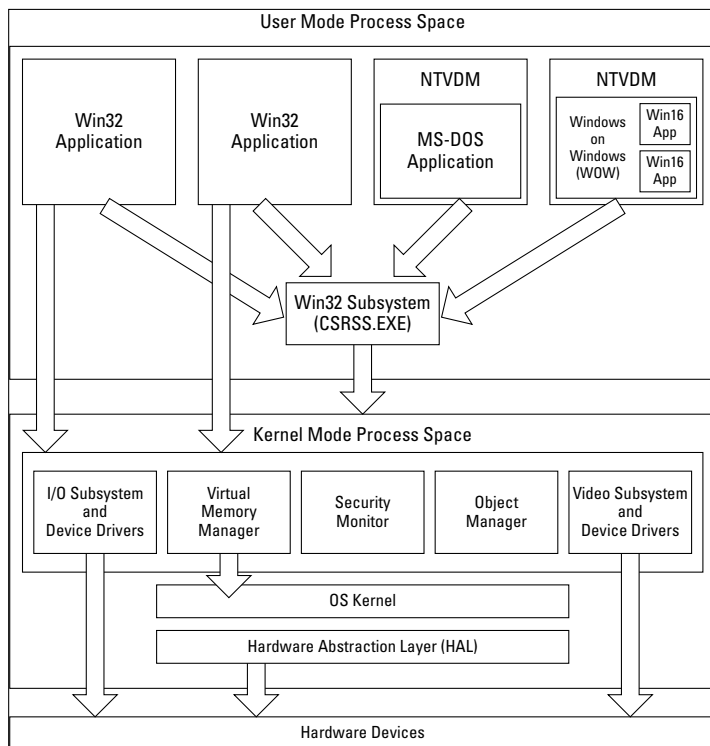


Figure 2-4: 32-bit applications run in their own execution area in Windows 2000 and Windows XP.

Address Space in the Windows NT–Based OS

The address space that the OS works with is 4GB of address space. This does not mean that you need 4GB of memory on the system to run the OS; you only need as much physical memory as your applications actual use to store items. Physical memory is only used when an application stores something into a block of address space.

Picture each block of address space being a label, and physical memory being a wall of mailboxes. If I store something in a mailbox, I attach

a label to that mailbox, representing what is stored there. When the application wants to store information, it may say, “Put this information in my address block 2045”; the OS takes that information and puts it into any free memory block, and then it keeps track of mapping between those two locations for the application. When the application wants that information back, it says, “Give me the information in address block 2045”, and the OS retrieves the information from where it was actually stored.

Running 16-Bit Windows Applications

The 16-bit Windows applications were originally designed to run under Windows 3.x. Some of these applications look for specific components in the operating system in order to function properly. Both Windows 9x and Windows XP emulate the 16-bit Windows environment to support 16-bit Windows applications. The actual way in which this emulation is performed represents one of the big differences between Windows 9x and Windows XP.

The 16-bit Windows applications are single-threaded applications that actually share a single unit of execution with the 16-bit Windows environment. This is what is referred to earlier in the chapter as *cooperative multitasking*. The entire 16-bit Windows environment executes via a single processor thread. This thread is then shared among all 16-bit Windows applications that are running in the 16-bit Windows environment.

The 16-bit Windows programming code is *non-reentrant*, which means that each section of code can be executed only once (or by one application) at any given time. When one application starts executing a section of code, it sets a flag — called the *Win16Mutex* — on that code. While the flag sets on the code, no other application is able to execute or enter that code. This is done to prevent multiple applications from executing the same section of code and thereby causing system-wide problems. One of the problems caused by the *Win16Mutex* flag is that any other program that wants to execute that section of code must wait until the current program finishes and the flag is removed. If the current application freezes, hangs, or crashes, it never removes its flag, thus the code is not released until the system is rebooted.



All 16-bit Windows applications suffer from these problems. Whenever possible, 16-bit applications should be replaced with equivalent 32-bit Windows applications to create an environment that is more stable and offers greater performance.

Executing in the Windows 9x environment

All 16-bit Windows applications run within a common memory space in the Windows 9x environment to reduce the load on system resources and to improve backward compatibility with applications that communicate with other applications through a shared memory address (a method of sharing data done in applications written prior to OLE [Object Linking and Embedding] and DDE [Dynamic Data Exchange]).

Having these applications run within one memory space causes issues when one application must be terminated from memory. Windows 9x is able to terminate only the *application proper* (the executable itself) from memory but is not able to remove any other components that were loaded into memory by the application. The 16-bit Windows environment never allowed you to track these “leftover,” lost resources, which could cause the application to hang or crash, so relaunching the application in the same 16-bit environment often resulted in the application and the environment becoming unstable. The best practice is to completely reload the 16-bit Windows environment, which, in Windows 9x, means rebooting the computer. (Rebooting is not necessary when a 32-bit Windows application crashes or hangs.)

All 16-bit Windows applications running under Windows 9x share a common message queue. The *message queue* contains keyboard and mouse input that has been directed toward the applications. If an application crashes or hangs while it has data in the message queue, the application does not remove its data from the queue. This causes the queue to become jammed and prevents other 16-bit Windows applications from getting their messages. When this happens, all data input to 16-bit Windows applications appears to be suspended. You notice this when you move your mouse over the top of any 16-bit Windows application — the mouse pointer remains an hourglass. To free up the message queue, you have to terminate the problem application.

If the stalled application has set the Win16Mutex flag on sections of code that the operating system requires to run its own 16-bit code, then the operating system may actually be prevented from running its own tasks. This can affect various parts of the GDI (Graphic Device Interface) and user interface (refer to the Windows 9x architectural diagram shown in Figure 2-3). The GDI is responsible for screen redraws and window movement, and the user interface is responsible for managing mouse and keyboard input. If the Win16Mutex flag blocks these two operating system components, your keyboard and screen may freeze up. In this situation, the only thing you can really do is reboot your computer.



The architecture of Windows 9x plays a big role in the OS's stability. This represents a big reason to upgrade to Windows XP, which has a dramatically different architecture.

Executing in the Windows 2000 and Windows XP environments

Running 16-bit Windows applications under Windows 2000, Windows XP, or Windows NT-based OSes is very similar to running them under Windows 9x — but *very similar* does not mean exactly the same. The areas of similarity include the following:

- ◆ **16-bit Windows applications run in a common memory space.**
- ◆ **16-bit Windows applications share a common message queue.**
- ◆ **16-bit Windows applications cooperatively multitask.**

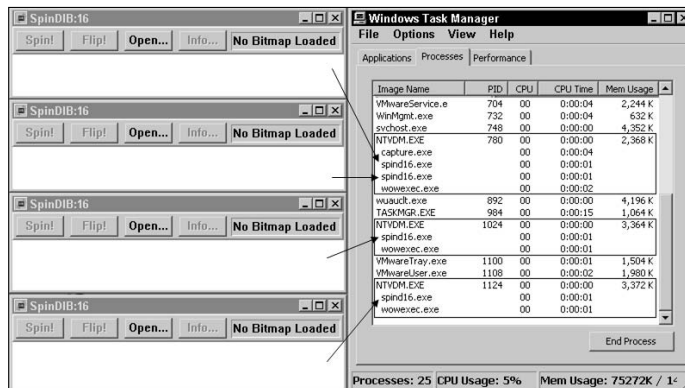
Even with these similarities, there are number of differences between Windows 2000, Windows XP, and Windows NT-based OSes and the Windows 9x environment. These differences start with where the application code is actually run. Within Windows NT-based OSes, all 16-bit Windows applications are executed from within a 32-bit Windows *NTVDM* (*virtual DOS machine*). This NTVDM is a complete emulated DOS environment onto which a complete emulated 16-bit Windows environment is loaded. It is within this environment that 16-bit Windows applications are executed. Although this seems like a small point, it means that the entire 16-bit Windows environment that is implemented within Windows XP is contained within a 32-bit environment. That makes the entire environment as safe and stable as any other 32-bit Windows application running on the system.

Within the emulated 16-bit Windows environment, all of the normal rules that apply to 16-bit Windows still apply, but they are only able to affect the other applications within the running NTVDM. Applications are still cooperatively multitasked. There is still a Win16Mutex flag, and code is still non-reentrant, but the reentrant code and other processes that can be halted are limited to this one environment. When you consider the separation of 16-bit applications and the fact that Windows XP does not contain any 16-bit code at the OS level, Windows XP has a major stability advantage over Windows 9x.

Another advantage to using the Windows 2000 and Windows XP environments is that 16-bit Windows applications can be run in separate memory spaces. When a 16-bit Windows application runs in a separate memory space, a new NTVDM is loaded along with a new *Windows on Windows (WOW) environment*, as shown in Figure 2-5. The benefits of running applications in separate memory space are preemptive multitasking between applications and process isolation. Although the application still cooperatively multitasks within this

Windows environment, there are no other applications for the clock to share cycles with. Since the entire 16-bit Windows environment is run within a 32-bit Windows process, the 32-bit Windows process preemptively multitasks other 32-bit Windows processes running on the system. With process isolation, each application can be run within its own 32-bit Windows process.

Figure 2-5: Task Manager shows that 16-bit applications can be loaded within separate Windows on Windows (WOW) environments.



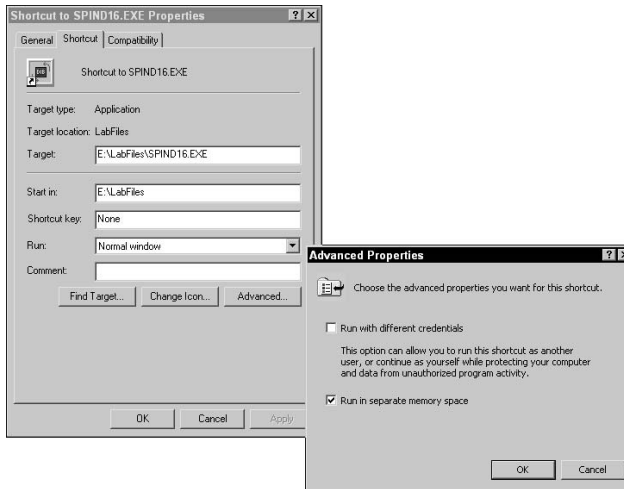
The benefit of separate processes is that when one 16-bit Windows application hangs or crashes, it does not affect any of the other applications running within other processes on the system. Any other 16-bit Windows applications that are running in the same process as the hung application are halted. Windows XP uses the default memory space to run all 16-bit Windows applications, with the exception of those that are run in separate memory spaces. A process running in a separate memory space is the only process that may run in that memory space.

There are several ways to start a 16-bit Windows application in a separate memory space:

- ◆ **At a command prompt by using** `start /separate <application name>`.
- ◆ **Create a shortcut to modify the properties on the Shortcut tab to run in a separate memory space, as shown in Figure 2-6.**
- ◆ **Modify the Open command for an associated file.**

Check out Book V, Chapter 4 for step-by-step instructions on altering the Open action for a file type.

Figure 2-6: Shortcuts represent one of the ways to open 16-bit applications in a separate memory space.



Any of these methods allow you to launch a 16-bit Windows application in a separate memory space, where it can run in isolation from other 16-bit Windows applications.

Terminating an application that is running under Windows XP can be accomplished through the Task Manager. Access the Task Manager by pressing Ctrl+Alt+Delete and clicking the Task Manager button or by right-clicking the taskbar and selecting Task Manager. The Task Manager has four tabs:

- ◆ **Applications:** Lists all running foreground applications
- ◆ **Processes:** Lists all running applications — foreground and background
- ◆ **Performance:** Provides performance statistics about your computer
- ◆ **Networking:** Provides performance statistics on your network interfaces

You can terminate an application from either the Applications or the Processes tab by selecting the application and clicking the End Task button. To terminate the entire 16-bit Windows environment, select the NTVDM that contains the application and click the End Task button.



You have seen the benefits of running applications in separate memory spaces, but you should also consider the drawbacks. There are two drawbacks to running applications in separate memory spaces:

- ◆ **Applications that communicate with each other through a shared memory address will not be able to communicate, since there will not be shared address.** Microsoft lists this as a minor problem because very

few applications that you will run actually perform this type of communication. (This has never been an issue with any of the applications I have worked with.)

- ◆ **You use additional system resources for every NTVDM and WOW that is loaded into memory.** The overhead associated with loading these resources is about half a megabyte for each NTVDM/WOW combination. This might not seem like a lot, but if you're running a large number of applications, it does add up.

You should now feel comfortable about how each operating system supports 16-bit Windows applications, and the options which you are able to control, modifying how the 16-bit applications function. As a CompTIA A+ Certified Professional, knowledge of these environments will allow you to address issues with applications and provide compelling reasons for late adopters to upgrade to the current Microsoft Windows OS.



Lab 2-1 gives you some practice examining the resources used by Windows on Windows (WOW). This lab requires that you obtain copies of specific 16-bit Windows applications from a copy of Windows 9x. These applications will be copied into the `c:\labfiles` folder on the lab computer. For this lab, you will have to use Windows NT 4.0, Windows 2000, or Windows XP.

Lab 2-2 will give you practice starting applications in their own memory space. This lab assumes that the required lab files have been installed to your disk, using the default installation path of `c:\labfiles`. If you have not completed Lab 2-1, then follow the steps at the beginning of that lab exercise to get the proper files in your `labfiles` directory. You can use Windows NT 4.0, Windows 2000, or Windows XP for these lab exercises.

Lab 2-1 and Lab 2-2 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Encountering incompatibilities

When running 16-bit Windows applications within Windows 9x or Windows XP, there is a chance that you will find incompatibilities with your operating system. Most developers want to make their applications run as fast as possible. This could mean that they follow nonstandard programming practices under Windows 3.x, which may actually cause problems when you attempt to run these applications under newer versions of Windows.

For example, if you have developed a fax application, you may use all standard Windows COM driver system calls in order to communicate with the modem. Upon investigating how the COM calls are actually made, you may decide that you can achieve more speed through the application by making

direct hardware calls. Direct hardware calls represent nonstandard communication with that COM device. This will not affect your application when running within the 16-bit Windows environment, but when running within a *simulated* 16-bit Windows environment, such as Windows 2000, these nonstandard calls might not be properly processed.

Part of the problem arises from the fact that in an attempt to provide better system stability, Windows 9x-based OSes and Windows NT-based OSes restrict direct access to some hardware components:

- ◆ Windows 9x-based OSes restrict access to COM devices.
- ◆ Windows NT-based OSes restrict access to most hardware components.



If an application requires direct hardware access to these components, it simply will not function under these newer versions of Windows, such as Windows XP. Microsoft has taken all possible steps to improve the ability to run older Windows-based applications, but will not take actions that will substantially reduce the stability of the OS.

Running MS-DOS-Based Applications

MS-DOS-based applications represent the oldest type of applications you are likely to run on your computer. MS-DOS applications are supported by Windows XP. MS-DOS-based applications have always been designed to run alone on a computer. They expect to be the only application ever running on that PC, and as such expect to see certain things, like no Windows presence. Microsoft uses virtual machines to provide a unique environment for these applications. These virtual machines simulate all the hardware that would normally be found in a computer, including

- ◆ **Keyboard**
- ◆ **Mouse**
- ◆ **Monitor**
- ◆ **COM ports**
- ◆ **RAM**
- ◆ **Disk drives**

With all of these components being virtualized, an MS-DOS-based application believes that it is running alone in a computer. The virtual computer's settings can be modified through a *program information file (PIF)*, which is discussed in the next section.

Program Information File (PIF) settings

Every MS-DOS-based application that executes on your computer starts by configuring a working environment from settings found in a *program information file (PIF)*. This is the case even if you are not aware that a .pif file is being used because a `_default.pif` file on your hard drive is used if no other .pif exists.

If you want to create a default .pif for a specific application then you should right-click the application and choose Properties. Even though there is not currently a .pif for the application, you will see all of the .pif related tabs, such as Screen, Memory, Font, and Misc. If you make any changes to the application settings on the tabs and select the OK button to close the dialog, then, a file named `<application>.pif` will be created in the same directory as the application executable. This .pif then becomes the default .pif for that application, since it is in the same directory as the application and is named `<application.pif>`. Unlike most files, the .pif extension stays hidden, even if displaying file extensions is enabled. You can tell that it is a .pif file if you view the properties of the file by right-clicking on it and selecting Properties. The Type of file will be Shortcut to MS-DOS Program, and the Font, Memory, Screen, and Misc tabs, are specific to .pif files, as shown in Figure 2-7. At any time after closing the dialog and the default .pif has been created, you can make changes to it by right-clicking on either the application or the .pif and getting properties, as either action will allow you to edit the settings. Even though you can access the settings through the properties of the application, you are really editing the .pif.


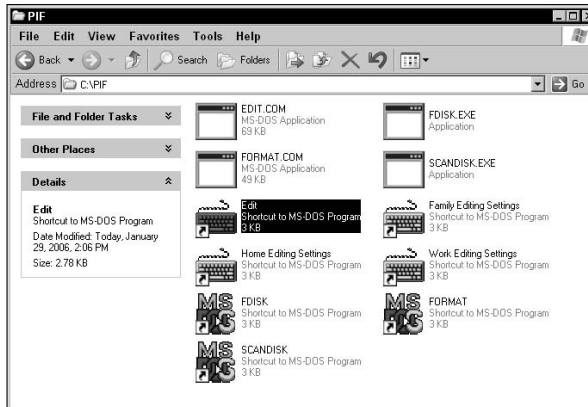


Figure 2-7: PIFs have several tabs that contain settings for a simulated MS-DOS environment.

You may also be confused by the fact that a `.pif` for an application looks exactly like a shortcut for a `.lnk` file. Figure 2-8 shows a series of `.pifs` that have been created for four MS-DOS-based applications. Notice that the first four icons in the window are MS-DOS application icons, and the last three icons are the default `.pifs` for the `fdisk.exe`, `format.com`, and `scandisk.exe` applications. In addition to the default `.pif`, `edit.com` has three additional `.pifs` that have been created for it, which are the keyboard-looking icons in the middle of the window. Each of these `edit.com` `.pifs` can have its own unique settings. To access the settings for a `.pif`, right-click it and select Properties. The Properties window of a `.pif` contains seven tabs, which are:

- ◆ **General**
- ◆ **Program**
- ◆ **Font**
- ◆ **Memory**
- ◆ **Screen**
- ◆ **Misc**
- ◆ **Summary**

Figure 2-8:
Each `.pif` for a common application can contain unique settings.



Each of these setting tabs is discussed in more detail in the sections that follow.

General tab

The first tab in the `.pif` Properties window is the General tab, and it is the same as the General tab for any other file on your hard drive. The General tab lists the Type of file you are viewing the properties for, which is a Shortcut to

MS-DOS Application. This tab will also list the location of the file, its size and size on disk, when it was created, modified, and accessed, as well as its attributes. You should also note the size of the .pif file — .pif files are very small (usually only 1K). A .pif can be differentiated from other shortcuts on your hard drive by its Type of file, since other files that have the shortcut arrow will only be identified as Shortcut, and not a Shortcut to MS-DOS Application.

Program tab

The Program tab, shown in Figure 2-9, lets you configure the following settings:

- ◆ **Cmd Line:** Specifies the path to the executable.
- ◆ **Working:** Specifies the working directory for the program; this is often used as the default save directory.
- ◆ **Batch File:** Specifies the name of a batch file or program that you would like to execute prior to launching the executable but after establishing the MS-DOS environment.
- ◆ **Shortcut key:** Specifies a keystroke that can be executed at any time to either launch or switch to this program.



With much trying, this author very rarely has seen the shortcut key in a .pif work the way it is supposed to. Rather than working, it does nothing.

- ◆ **Run:** Specifies whether the program is supposed to run in a normal, maximized, or minimized window.
- ◆ **Close on Exit:** Specifies whether the application should close its window when it completes execution. Note that having the program close on completion sometimes hides error messages that could be useful for troubleshooting.

There are also two buttons at the bottom of the Program tab. One is Change Icon, which allows you to change the icon that is used by the .pif. The other is Advanced, which allows you to modify some additional settings.

Changing the icon is purely cosmetic, but you can specify an icon from any Windows-based executable or from a few of the system DLLs. Every Windows-based executable contains an icon list (with the first icon in the list having an index number of 0) that includes the executable's icon (iconnumber 0) and icons for its documents or files. Icons can also be found in several system DLLs, such as `moricons.dll`, `pifmgr.dll`, and `shell32.dll`. `shell32.dll` contains many of the default system icons that Windows uses.

Clicking the Advanced button brings up the Windows PIF Settings dialog box, which enables you to override the default `autoexec.bat` and `config.sys` files. The default files — `autoexec.nt` and `config.nt` — are found in the

`%SystemRoot%\system32` directory. These two files are used to create the default MS-DOS environment for Windows NT-based OSes. You can create new files with any names and specify the files here. There is also a check box to enable compatible timer hardware emulation, which reduces the rate at which the computer sends timing signals to the application. In short, it makes the application think that it is running on a slower computer.

Figure 2-9:
The Program tab allows you to specify information related to the executable file.



If you want to know how to modify your `autoexec.nt` and `config.nt` files, you should read through the `autoexec.bat` and `config.sys` sections of Book V, Chapter 6.



In Lab 2-3, you create a `.pif` to launch the MS-DOS `edit` command. You can perform this lab using Windows 9x, Windows 2000, or Windows XP. Lab 2-3 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Font tab

You use the Font tab (shown in Figure 2-10) to specify the type(s) of screen font(s) that can be used by the window that contains the running MS-DOS application. These fonts can be either TrueType fonts or bitmap fonts. TrueType fonts are created on the fly by using a formula, and thereby can be scaled to any size, while bitmap fonts are drawn with pixels or squares to a specific size. Bitmap fonts are faster to draw, but do not scale. By default, Windows uses both and chooses the best font based on the window size.

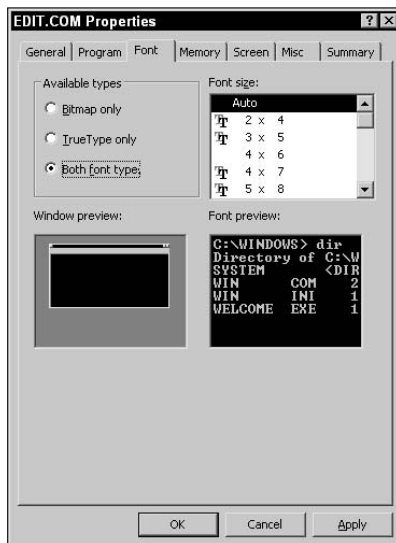


Figure 2-10: Fonts are chosen for readability of output.

Memory tab

The Memory tab enables you to specify limits for application memory. As seen in Figure 2-11, you can configure four types of memory:

- ◆ **Conventional memory**
- ◆ **Expanded (EMS) memory**
- ◆ **Extended (XMS) memory**
- ◆ **MS-DOS protected mode (DPMI) memory**

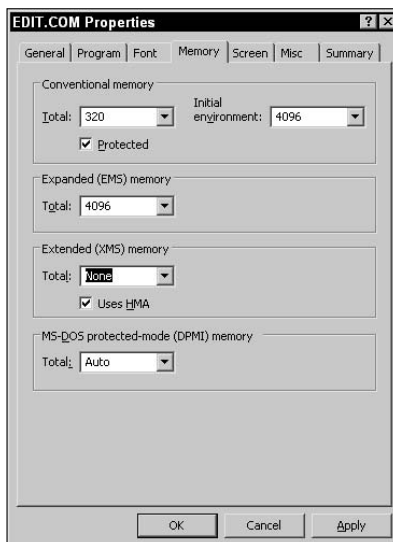


The default setting for each of these types of memory is either Auto or None. For more information about the different types of memory and how to configure them, consult Book V, Chapter 6.

For MS-DOS applications, *Conventional memory* is where the application actually operates. When MS-DOS first hit the market, computers only made use of the first 640K of memory that was installed, with the rest of the 1024K being used to access hardware components and for drivers to load. Windows applications, which came out later, were designed not to use this area, allowing MS-DOS to continue to use the Conventional memory space. Since an entire computer is emulated for MS-DOS applications, this first 640K is also emulated, and although the application thinks that it is functioning in the first 640K of memory, Windows could be using any space that the Virtual

Memory Manager has available. On a normal computer, some of the conventional memory is used to load drivers and TSR (Terminate Stay Resident) applications. A TSR would be similar to service in Windows, since it runs in the background. When it is time for an application to run, it may not have all of the first 640K of memory available to it. The Total Memory drop-down menu allows you to specify the amount of Conventional Memory that is available to the application. When the Total Memory drop-down menu is set to Auto, the application is provided with as much conventional memory as it requests. It seems odd that more conscientious programs actually have problems with this setting. The conscientious programs query the operating system before launching to inquire about the amount of free memory available, but rather than returning a value, Windows asks the question, “How much memory do you want?” This merely confuses the application, which then generates an error message stating that an insufficient amount of memory exists to run the program. If that happens to you, you can adjust the amount of conventional memory that is available to the application. This drop-down menu is measured in kilobytes.

Figure 2-11: Memory that is available to the application is emulated through system memory.



You can allocate up to 640K of conventional memory to the application, but you should allocate only what the application truly requires. If your application requires the entire 640K and you are still having problems, you can allocate up to an additional 4K of memory to hold environment variables. This memory is added by making use of the *Initial Environment* drop-down menu, which is measured in bytes.

There is also an option to *protect* the conventional memory, which prevents the memory space from being swapped out to the Windows swap or paging file. (Some applications crash when they are swapped out to the swap or paging file.)



Protecting memory should be avoided unless absolutely necessary because it locks the application into physical RAM, reducing the amount of physical memory that is available for other Windows-based applications.



When dealing with the virtual memory file that is found on your hard drive, Windows 9x refers to this file as the *swap file*, while Windows NT-based OSes refer to it as the *paging file*.

When 640K of memory became too restrictive for developers, hardware vendors developed expansion memory that was able to be added to a computer by using ISA expansion slots, and using a small portion of the memory between 640K and 1024K as a swap area. The application would request certain portions of memory, which would then be swapped from the Expanded memory card into conventional memory. This memory was called *Expanded (EMS) memory*. In order to use expanded memory under Windows 9x-based OSes, you have to load the Expanded Memory Manager (`emmm386.exe`) in your `config.sys` file, which allocates the specific memory address blocks that will be used for the swapping process. You will find `emmm386.exe` in your Windows directory, and you can review Book V, Chapter 6 to learn how to work with your `config.sys` file. After the Expanded Memory Manager is loaded, you can specify up to 16MB of expanded memory to be available to the application by using the Expanded Memory drop-down menu, which is measured in kilobytes.

As computers started to be produced with more than 1MB of memory, people wanted to make use of this extra memory above the 1MB mark, but MS-DOS natively was only able to access 640K of memory. *Extended (XMS) memory* is the memory that was above 1MB. Most MS-DOS applications do not use extended memory, since it could only be accessed by Protected Mode applications or operating systems like Windows. Extended memory is sometimes also called Protected Mode memory, since it is usually accessed by Protected Mode applications. If you load a driver like `himem.sys`, then you can access this extended memory. You can allocate up to 16MB of extended memory for the application using the Extended memory drop-down menu which is measured in kilobytes. The MS-DOS environment can use the high memory area that is provided by extended memory by specifying an option when loading `himem.sys` in the `config.sys` file. The high memory area is the first 64K of Extended memory and through a design glitch, can be accessed directly by MS-DOS to load drivers.

Very few applications use *DOS protected-mode interface (DPMI) memory* these days, and most of those are games like DOOM and Descent. This memory management technique allows MS-DOS-based applications to use extended memory through special MS-DOS and BIOS calls. The big advantage that is provided by DPMI memory use is that it allows an MS-DOS application direct memory access over 1MB, which provides better performance than using extended memory (which uses swapping to move information between extended memory and conventional memory areas). If your application does make use of DPMI, you can specify up to 16MB of DPMI memory. You allocate DPMI memory specifying the amount in the DPMI drop-down menu, which is measured in kilobytes.

Screen tab

The Screen tab enables you to specify whether the application is to run full-screen or in a window. You can also set the number of lines of text that will appear on a screen; the default is 25 lines, but you can specify up to 50 lines. The benefit of adjusting the number of lines on the screen is that you can display more data at any given time, but many applications don't accept this setting. If you are running the application within a window, then you can specify whether the toolbar is visible and whether Windows should automatically restore your previous settings upon the next startup of the application. Figure 2-12 shows these settings and the other display settings that can be modified on the Screen tab.

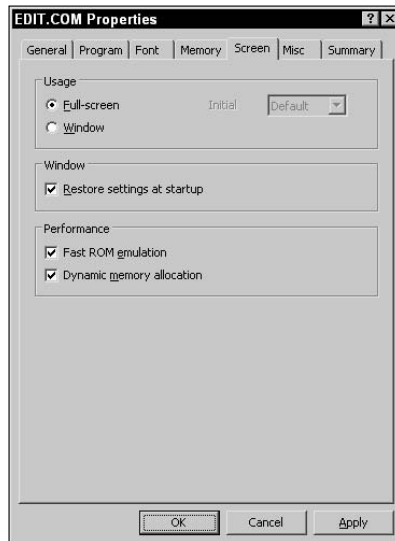


Figure 2-12: Screen settings hold several screen- and video-related settings.

There are also two check boxes that play a big role with video performance. The first of these is Fast ROM Emulation, which copies the contents of the video ROMs on your system into RAM to optimize performance. Fast ROM emulation is incompatible with some applications. The second check box enables Dynamic Memory Allocation, which allocates only a text mode video buffer for the application. Some applications work in graphics mode from time to time, and when they do, if this option is activated, Windows dynamically changes the size of the buffer to support graphics mode. Once again, this may cause some applications to crash. When you are troubleshooting problems with an MS-DOS-based application, it's worth checking out these settings. You don't necessarily have to disable the settings, but if you're having video-related problems, there is a good chance that the settings should be disabled.

Misc tab

The Misc (Miscellaneous) tab is a grab bag of settings that affect how the MS-DOS application behaves. These settings are shown in Figure 2-13. The settings affect how the program operates in the foreground or background, how it terminates, the number of clock cycles it receives, and the shortcut keys that it uses.

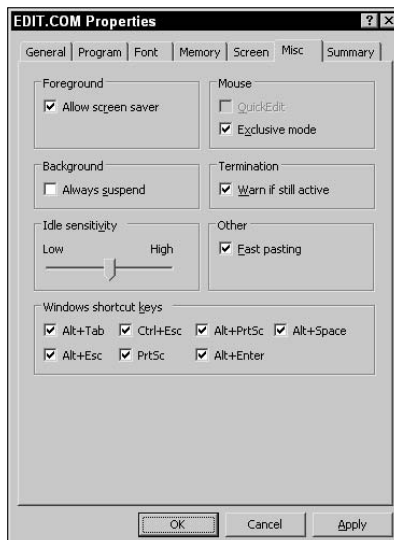


Figure 2-13: The leftover settings are all stored on the Misc tab.

The first of these settings, Foreground, lets you choose whether the Windows screen saver is allowed to function when the application is active. If your application is running in full-screen mode, it has to be minimized in

order for the Windows screen saver to be active. If your application does not like running from within a window, or while minimized, then this setting could cause problems. Not allowing the Windows screen saver to be active leaves your screen running 100 percent of the time when the application is active in the foreground. This may not be such a big deal, though, because many monitors now have an energy-saving standby feature.

The Background settings specify whether the application is suspended when it is the background task, which means it receives clock cycles only when it is the foreground application. If the application requires clock cycles at all times, having it suspended in the background may cause it to crash. Most applications don't like being suspended when in the background, although if you have an application that uses a large number of clock cycles, you may want to try this option.

The Termination setting of Warn If Still Active applies to MS-DOS windows that are closed by using the Close box in the top-right corner of the window. This warning setting has been designed to help prevent data loss when the application is improperly closed by displaying a confirmation dialog box that allows you to cancel the closing of the application. When you try to close normal Windows applications, they usually ask you whether you want to save your changes, but closing an MS-DOS-based application by clicking the Close box bypasses the normal Close functions and promptly terminates the MS-DOS session and therefore the application.



You should always exit your MS-DOS application from within the program, using whatever commands or options are required. Unlike Windows programs that have a standard Exit in the File menu, MS-DOS applications usually have their own special commands for exiting or closing the application, such as key combinations using CTRL or ALT, and Q, X, or W.

The Mouse settings of Quick Edit and Exclusive Mode affect the way the mouse functions for the MS-DOS-based application. Quick Edit allows you to use Windows-like mouse movements to select, delete, and insert text. Quick Edit applies only to Windows 9x computers, not to Windows XP computers. Exclusive Mode may seem a little odd; when enabled, it locks your mouse into the confines of the MS-DOS-based application. This is an option because some MS-DOS-based applications have nonstandard mouse routines. Normally when your application is running in a window, you can move your mouse out of the window at any position around the window and bring it back in at any other position.

However, in some applications that use nonstandard mouse routines, the application thinks the mouse is still in the exit location while Windows thinks the mouse is at the re-entry location, which leaves your application mouse pointer and your Windows mouse pointer in different locations. With most programs that use standard mouse routines, when your mouse is

brought back into the application, the location of your pointer is resynchronized. If your mouse is not synchronized in your application, then set your mouse in Exclusive Mode. You can still move your mouse outside of the window by using any of the standard Windows shortcut keys, such as Alt+Tab. When you click on the application to make it active again, Windows synchronizes the mouse pointers for you.

The Fast Pasting option uses Windows-based routines, rather than standard MS-DOS routines, for inserting text into your application. This works with most MS-DOS-based applications but causes problems for some. You know that you have a problem when pasting operations cause errors or disable the application, so disable Fast Pasting in these cases.

The Idle Sensitivity setting specifies how your application responds to having its clock cycles reduced. When multiple applications are running on your system, Windows splits clock cycles among all applications that are actively requesting CPU time. When it comes to MS-DOS-based applications, Windows provides clock cycles only to applications that appear to require them. The problem is that some applications that require CPU time do not appear to need it. If an application seems to stall or crash for no apparent reason, you may want to lower its sensitivity to being idle. Selecting this option means that the application receives additional clock cycles even when Windows doesn't perceive that it needs them.

The last section on the Misc tab provides a list of standard Windows shortcut keys. This is different from the Shortcut Key on the Program tab, which is used to launch the .pif. If your application uses any of the shortcut keys, you must clear the appropriate check boxes here to prevent the standard Windows-based responses to the shortcut keys. For example, if your application uses the Alt+Tab shortcut, and you have not cleared the Alt+Tab check box on this tab, every time you press Alt+Tab, you will be switched to another application. Unless you clear the check box here, the Windows shortcut keys always take precedence.

Summary tab

The Summary tab contains summary information about the file. This is a standard tab for all files on Windows 2000 and newer OSes. These properties will be used by the Index Service and can help you locate files on your hard drive. The fields on this tab have no effect on how the .pif can be utilized.

Understanding incompatibilities

When you examined 16-bit Windows-based applications earlier in this chapter, you looked at potential incompatibilities that could arise from direct hardware access. The same incompatibilities apply to MS-DOS-based

applications. If your application attempts to make direct hardware calls instead of standard MS-DOS-based system calls, it will not function under Windows 9x or Windows XP. In an effort to be more backward compatible, Windows 9x places restrictions on fewer hardware devices, but Windows XP is far less forgiving. Basically, if you have an application that makes direct hardware calls, it will not work under the newer Windows OSes.

Windows XP Compatibility Modes

Many of the applications written for DOS or earlier version of Windows will work with Windows XP with a few minor tweaks. Starting with Windows XP, Microsoft has provided a new set of options to aid with backward compatibility for applications that are not included with the operating system. These features are found on the Compatibility tab in the Properties dialog box of the executable, which is shown for an application and a .pif in Figure 2-14.

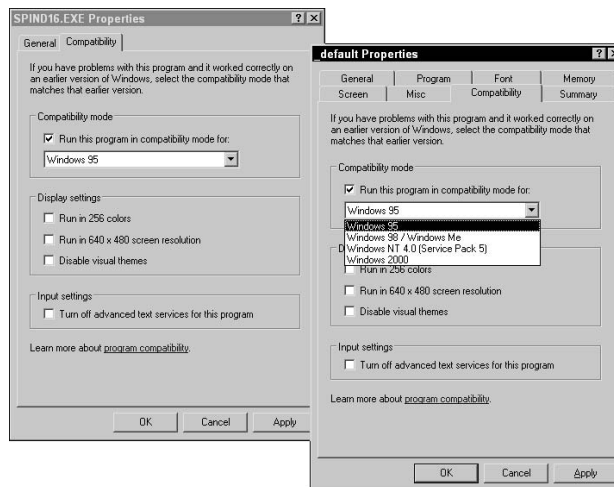


Figure 2-14:
The Compatibility settings for an application and a .pif.



These options have been designed to be used with locally installed applications. You can use them with applications from CD-ROMs or from a network drive, but the settings will not be saved.

Since many applications have been written to take advantage of specific features found in various versions of the operating system, the Compatibility Mode section allows you to emulate the following OSes:

- ◆ **Windows 95**
- ◆ **Windows 98/Windows ME**

- ◆ **Windows NT 4.0 (Service Pack 5)**
- ◆ **Windows 2000**

This does not emulate the entire OS, but rather disables some newer APIs (Application Programming Interfaces) and emulates some older ones, thereby improving compatibility.

Another common problem revolves around screen settings. Many times, I have had to change my screen resolution or color depth to run a specific application (and then reset it later). The Display Settings section allows you to automatically change the display settings for the application to one of the following settings:

- ◆ **Run in 256 colors**
- ◆ **Run in 640x 480screen resolution**
- ◆ **Disable visual themes**



If you have an application that has problems displaying menus or toolbar buttons, you can attempt to use the Disable Visual Themes setting. This removes the modern rendering of buttons, with the Windows XP smoothness, drop shadows, and so on.



Finally, under Input Settings, you have the option to Turn Off Advanced Text Services for This Program, which disables the features of Windows XP that do speech-to-text and handwriting recognition. The additional APIs running to support those services can cause unexpected input into your application.



Most of the settings that can be adjusted or disabled on the Compatibility tab can also be changed elsewhere, such as the Performance Options, on the Advanced tab of the System Control Panel, which will change the setting globally for the computer.

If you have an old application that will not run properly on Windows XP, do not forget about the Application Compatibility Settings.

Application Install, Start, and Load Errors

Many of Windows XP's application errors are related to permissions to the file system when using NTFS and the Registry. Windows XP has default permissions that restrict changes to files, addition of files to specific directories, and changes to portions of the Registry. In addition to these restrictions, many applications that attempt to directly access hardware will fail. If the

application is a modern 32-bit Windows application that is supposed to work with Windows XP, then most installation errors can be traced to not meeting either the minimum installation requirements or permissions.

Microsoft has made these permissions changes to prevent most users from being able to install their own applications and make unauthorized changes to the operating system. One way to see whether the problem is with a user installing or launching an application is to temporarily add the user to the local Administrators group and see whether the problem is resolved. If that resolves the issue, then you know that the problem is related to permissions to one of these locations.

If you can't get permissions information from the application vendor or can't grant the user elevated group membership, the SysInternals (www.sysinternals.com) Web site has utilities such as RegMon and FileMon that can help you identify the permissions changes required.

Getting an A+

This chapter examines application management from installation to execution. The following points are covered:

- ◆ 16-bit, 32-bit, and DOS-based applications function under Windows 9x and Windows XP.
- ◆ There are core architectural differences between Windows 9x and Windows XP. Support for these types of applications differs based on the architectural differences.
- ◆ System stability and application support seem to go hand in hand for these operating systems, and Windows XP provides more stability but a lower level of application support than Windows 9x.
- ◆ The new application compatibility features in Windows XP allow it to support a wider range of older Windows applications.

Prep Test

- 1** Which of the following types of multitasking are supported by Windows XP? Choose all that apply.
 - A Share-based multitasking
 - B Equality multitasking
 - C Cooperative multitasking
 - D Preemptive multitasking

- 2** Which of the following types of multitasking provides better sharing of CPU time between processes?
 - A Share-based multitasking
 - B Equality multitasking
 - C Cooperative multitasking
 - D Preemptive multitasking

- 3** Which operating system allows for 16-bit Windows applications to be preemptively multitasked?
 - A Windows XP
 - B Windows 98
 - C Windows 95
 - D MS-DOS

- 4** Which of the following are benefits of 32-bit Windows applications? Choose all that apply.
 - A They share a common message queue.
 - B They run in separate memory spaces.
 - C They are preemptively multitasked.
 - D They support running multiple applications in one virtual machine.

- 5** How large is the memory space that each 32-bit Windows application has to work with?
 - A 16MB
 - B 1GB
 - C 2GB
 - D 4GB

- 6** Which of the following operating systems support 16-bit Windows applications? Choose all that apply.
- A Windows XP
 - B Windows 98
 - C Windows 3.1
 - D MS-DOS
- 7** How can you preemptively multitask 16-bit Windows applications when using Windows XP?
- A Right-click the application icon and choose Launch in Separate Memory Space.
 - B Choose Start→Run, use the Browse button to locate the application, and then select the Start in Separate Memory Space check box.
 - C Type `start /separate <application>` in a command prompt.
 - D This cannot be done.
- 8** What effect does a 32-bit Windows application have on the rest of a Windows XP system when it hangs or crashes?
- A It halts 32-bit Windows applications that are running on the system.
 - B It halts 16-bit Windows applications that are running on the system.
 - C It halts MS-DOS-based applications running on the system.
 - D It has no effect on other applications.
- 9** What effect does a 16-bit Windows application have on the rest of a Windows XP system when it hangs or crashes?
- A It halts 32-bit Windows applications that are running on the system.
 - B It halts 16-bit Windows applications that are running on the system.
 - C It halts MS-DOS based applications running on the system.
 - D It has no effect on other applications.
- 10** What extension is assigned to shortcuts to MS-DOS applications?
- A .pif
 - B .lnk
 - C .dos
 - D .bin

- 11** What is the name of the default `config.sys` file that is processed by Windows XP when launching an MS-DOS-based application?
- A `config.sys`
 - B `config.nt`
 - C `ntconfig.sys`
 - D `ntconf.sys`
- 12** Bill has modified the settings for his application `c:\customapp\custom.exe` by editing `c:\customapp\custom.pif`. When he launches the application using the shortcut that is on his desktop, he does not get the changes that he made. Why?
- A By using the `.pif` on his desktop, he has requested the settings stored in that `.pif` rather than those in `c:\customapp\custom.pif`.
 - B He probably clicked the Cancel button rather than the OK button when leaving the Settings window for the `.exe`.
 - C Application setting changes for `.pif` files require you to reboot before they are in effect.
 - D `.pif` files do not affect the execution of programs.
- 13** How many applications can be assigned to an NTVDM when using separate memory spaces to launch the applications?
- A 1
 - B 2
 - C 4
 - D 8

Answers

- 1 C, D. Windows XP supports cooperative and preemptive multitasking. See *"Getting the Most Out of Multitasking."*
- 2 D. Preemptive multitasking provides more even distribution of CPU time between applications. Review *"Preemptive multitasking."*
- 3 A. 16-bit applications can be preemptively multitasked within Windows XP by running each application in a separate memory space. Check out *"Running 16-Bit Windows Applications."*
- 4 B, C. 32-bit applications use separate message queues so that when one application hangs, it does not affect the others. Only 16-bit Windows applications may be run in a common VM. Running applications in separate memory spaces and preemptive multitasking are benefits of running 32-bit applications. Peruse *"Benefiting from 32-bit applications."*
- 5 D. Each application or VM runs in a 4GB memory space. Take a look at *"Running 32-Bit Windows Applications."*
- 6 A, B, C. Only Windows-based OSes can run Windows-based applications. Windows 3.1 was only capable of running 16-bit applications, and Windows XP and Windows 98 support 16-bit applications in an effort to be backward compatible. Peek at *"Running 16-Bit Windows Applications."*
- 7 C. Using the `start` command allows you to run applications in separate memory spaces. Look over *"Executing in the Windows 2000 and Windows XP environments."*
- 8 D. When a 32-bit Windows application crashes, it has no effect on other applications running on the system. Study *"Executing in the Windows 2000 and Windows XP environments."*
- 9 B. When a 16-bit Windows application crashes, it halts other 16-bit Windows applications in the same memory space. MS-DOS-based applications and 32-bit Windows applications should continue to operate. Refer to *"Executing in the Windows 2000 and Windows XP environments."*
- 10 A. `.pif` is used as the extension for shortcuts to MS-DOS-based applications. Examine *"Program Information File (PIF) settings."*
- 11 B. `config.nt` is the default `config.sys` file that is processed by MS-DOS `.pif` files. See *"Program tab."*
- 12 A. Application settings are used from the launching `.pif`. If the application is launched directly, the `.pif` in the application's directory with the same name as the application is used; if there is no matching `.pif` in the directory, the default system `.pif` is used. Review *"Program Information File (PIF) settings."*
- 13 A. When starting applications in separate memory spaces, you can assign only one application to each NTVDM. Peruse *"Executing in the Windows 2000 and Windows XP environments."*

Chapter 3: Optimizing the Windows Environment

Exam Objectives

- ✓ **Optimizing operating systems by modifying virtual memory settings**
- ✓ **Using utilities to monitor and identify optimization areas**
- ✓ **Optimizing hard drives and temporary files**
- ✓ **Managing services and startup of applications to optimize system performance**

Few things in life are perfect, and if you put them on a scale from perfect to lousy, then most things would fall somewhere in the middle. Computers play a part in my life, and I can safely say that they are not perfect, especially when dealing with speed and performance. I regularly work with other people's computers and find the responsiveness of many computers far less, very far less, than perfect. Over time, if left to its own devices, your computer will slow down, sliding down the scale from the perfect end to the lousy end. When I comment on the slow response of a computer that just completed a 15-minute boot and logon, I am often surprised, or absolutely floored, when the owner says that it is acceptable. Although some people find *acceptable* to be, well, acceptable, as the saying goes, *good enough* is never good enough. As a CompTIA A+ Certified Professional, you should be able to identify areas that may be causing a slow down and resolve the issues.

In this chapter, you look at the major problems that are responsible for many slowdowns and how to avoid them.

Identifying Areas of System Bottlenecks

The neck of a bottle, which has a substantially smaller diameter than the rest of the bottle, restricts the flow of liquid from the bottle. Bottlenecks occur on your computer as well. They happen when most system resources are fine, except for that one sub-system that is heavily overused. In order to perform any task, a limited number of resources can be applied to a task, and when the task exceeds the available resources, you run into problems. Computers have four critical resources that may be the source of a bottleneck:

- ◆ **Processor**
- ◆ **Memory**
- ◆ **Disk**
- ◆ **Network**



All system bottlenecks will occur in one of these four resources, which are the main sub-systems in a computer.

In addition to the overuse of resources, hardware errors might also cause problems for you. This chapter does not examine the possible hardware errors that can cause problems or issues that may be specific to a particular type of application or service. Most hardware issues will be dealt with in Book IV, Chapter 2, which is devoted to troubleshooting.

Using Monitoring Tools

In order to effectively diagnose a problem in one of the four critical resources (processor, memory, disk, and network), you need a way to monitor what is going on with your system. In the following sections, you look at two tools that you can use to diagnose problems in these resource areas.

Task Manager


Task Manager or `taskman.exe` is a nice, quick, and simple tool. It is not as full-featured as Performance Monitor (see the following section), but what it lacks in features, it makes up for in simplicity and speed. To open Task Manager, you can press `Ctrl+Alt+Del` and click the Task Manager button in the Security dialog box. If you're using Windows XP Home, then just pressing `Ctrl+Alt+Del` opens Task Manager. You can also open Task Manager by right-clicking an empty area of the Taskbar and choosing Task Manager from the context menu.

Performance tab

You can use Task Manager to diagnose processor, memory, and (with Windows XP) network bottlenecks. Figure 3-1 shows the Performance tab of Task Manager, which is typically the tab you will use first to identify a problem.

The Performance tab of Task Manager provides critical data about your system's performance. This tab is broken down into two main sections: the graphs and the numeric data. In the graph section, you see CPU usage and PF (Page File) usage information. In the numeric section, you see information on processes and memory usage. At the very bottom of the window, you see summary information in the status bar.

Figure 3-1: The Performance tab of Task Manager is usually the first place to look for problems.



You might notice that the Task Manager window stays on top of all others, which is normally preferred; however, if you want it to go to the background, then choose Options→Always on Top to deselect this option.

In the graphs section, you see the graph for the overall CPU usage, which is the same image shown down in the System Tray next to the clock (this image appears automatically when you open Task Manager). *CPU usage* is an average of CPU usage over all of the processors in your computer.

Next to this graph is a line graph showing you historic CPU usage over the last couple of minutes. If you have multiple processors, processors with hyperthreading, or dual core processors, then you will see multiple line graphs, each in its own small window. If the graph is too small, you can resize Task Manager. If you have multiple graphs, you can get a single graph showing the averages by choosing View→CPU History→One Graph, All CPUs. If you choose View→Show Kernel Times, you get a red line showing kernel processor utilization to help identify whether the problem is related to kernel processes or user processes. It is not uncommon for processor utilization to jump to 100%, but if it is consistently above 80% or 90%, you're likely experiencing a slowdown.

Page file usage is also recorded in the graphs section. *Page file usage* is actually application-specific virtual memory usage, which includes paging file usage and a portion of the physical memory that the application is using. Just like CPU utilization, there are two graphs here, a histogram for instantaneous usage and a line graph showing the historical usage. If you are running low on space in your page file, then you are likely running short on memory.

In the numeric section, you find totals for handles, threads, and processes. Think of *processes* as applications, but not all applications run in windows you can see — some, like services, run in the background, and these are included in the total number of running processes. Each process is composed of *threads* of code that are executed. Old applications from the 16-bit Windows and MS-DOS days are *single threaded*, which means that the program runs a single thread of code from beginning to end. Newer applications are *multi-threaded* and can run several different threads of code, each performing a different task, all working together to perform the task that processor is attempting to accomplish. On a multi-processor computer, these threads can execute simultaneously, improving the overall speed of the process.

Handles are resources that the processes are using, such as memory locations, files, or registry keys. Most likely you will be primarily concerned with the number of processes that are running, because each running process consumes some of the total processor cycles that are available.

The other three numeric sections all deal with memory — physical memory, kernel memory, and commit charge. *Physical memory* is the actual RAM that is installed in your computer. The physical memory section lets you see how much physical memory is on your computer, the amount that is currently free (or available), and the amount that has been allocated to system cache or disk caching for the disk drives. The cache space dynamically adjusts as demands for memory go up, but the goal is to improve disk access by caching in fast memory.

Kernel memory is memory that is allotted to the kernel. The kernel memory section lets you see how much memory has been allocated to the kernel, which is in charge of running all operations in your computer. If a driver has a memory leak, then you may see kernel memory increase higher than normal for your computer, which is typically below 100MB — the short-term solution is to reboot your computer.

Kernel memory is split between paged and non-paged memory. *Paged memory* is part of the *page pool*, which is the memory that is swapped between physical RAM and the page file on the hard drive.


Commit charge is the amount of memory that is in use, or committed. This section shows you current usage, total available (including the page file), and the peak usage since the last reboot.

Networking tab

In most cases, looking at the data on the Performance tab leads you to either the Processes tab or the Networking tab. On the Networking tab (shown in Figure 3-2), you see a line graph with lines for each network interface you

have on your computer, representing the percentage of bandwidth that is being used by each one. If you would like to see additional information on your graph, you can choose one of the options on the View⇨Network Adapter History sub-menu. This allows you to add Bytes Sent and Bytes Received to the existing Bytes Total line on the graph. This is useful for seeing whether a network-related problem is due to data coming in or going out.

Figure 3-2:
The Networking tab of Task Manager can help you identify network use problems.



You can also find the same information from the graph presented numerically at the bottom of the window. If you would like to see more data, choose View⇨Select Columns. This opens the tool to a lot more troubleshooting capabilities by allowing analysis of many network-related counters, such as the breakdown of bytes sent, bytes received, and traffic type being unicast or nonunicast. *Unicast* network traffic is information that is sent only to your computer, while *nonunicast* traffic is simultaneously sent to multiple computers on your network.

Processes tab

If the problems are not network-related, then you may be led to look at the Processes tab, which is shown in Figure 3-3. This tab lists all of the running processes on your system in columns. These columns can be sorted in ascending order by clicking the column heading, or descending order by clicking the column heading a second time. The default columns include Image Name, User Name (for the user who is executing the process), CPU (for the percentage utilization of this process), and Memory (for the amount

of memory that this process is using). If you are logged on as an administrator, you can also choose to Show Processes from All Users to see processes other than your own.

Figure 3-3:
The
Processes
tab of Task
Manager
helps
identify CPU
utilization
problems.

Image Name	User Name	CPU	Mem Usage
System Idle Process	SYSTEM	88	28 K
taskmgr.exe	Bob	05	2,360 K
wowexec.exe	Bob	00	
capture.exe	Bob	00	
ntvdm.exe	Bob	02	10,764 K
explorer.exe	Bob	02	11,536 K
winlogon.exe	SYSTEM	02	1,440 K
System	SYSTEM	02	60 K
wuauclt.exe	SYSTEM	00	6,344 K
VMwareService.exe	SYSTEM	00	736 K
wuauclt.exe	Bob	00	756 K
inetinfo.exe	SYSTEM	00	1,588 K
cisvc.exe	SYSTEM	00	528 K
cid daemon.exe	SYSTEM	00	188 K
spoolsv.exe	SYSTEM	00	1,360 K
WINWORD.EXE	Bob	00	3,548 K
msmsgs.exe	Bob	00	1,580 K
VMwareUser.exe	Bob	00	748 K
VMwareTray.exe	Bob	00	508 K
csrss.exe	LOCAL SERVICE	00	1,426 K

Show processes from all users End Process

Processes: 33 CPU Usage: 20% Commit Charge: 157880K

In most cases, you will be able to locate the process that is using up most of the CPU cycles and slowing your computer down. If you need to see additional troubleshooting information, you can choose View→Select Columns. This lets you select from many other counters, such as Virtual Memory Size, Page Faults, and Peak Memory Usage, to list just a few.

Page faults occur when information that your computer wants is not in physical memory and must be initially loaded into memory or read from the page file on the hard drive. If the number of page faults takes a sharp rise, you may also notice that your available memory is low.

If you right-click any running task on the Processes tab, you will have several options, including ending the process, ending the process tree, or changing the application priority. If you end the process tree, you terminate the process and any other processes that it started. Changing a process's priority changes the percentage of CPU time that an application gets when it is running.



Be careful when using high and real-time priorities; they can cause system instability. If you have a highly active program, such as a counter, then setting it to real time priority can allow it access to all CPU cycles on your computer, which may cause the system to ignore user input from the mouse or keyboard, preventing you from doing anything else with your computer.

Applications tab

The last tab to look at is the Applications tab. This tab shows you what applications are running in the foreground (as shown in Figure 3-4). These applications should list Running in the Status tab. If the application is not listed as running, it may not be responding. At this point, you can either give the application more time to respond — it might just be busy — or you can end a task by selecting the application and clicking the End Task button. If you would like to bring an application's window to the foreground, then you can select the application and click the Switch To button. If you need to launch a new application, then you can click the New Task button, which brings up the Run dialog box.

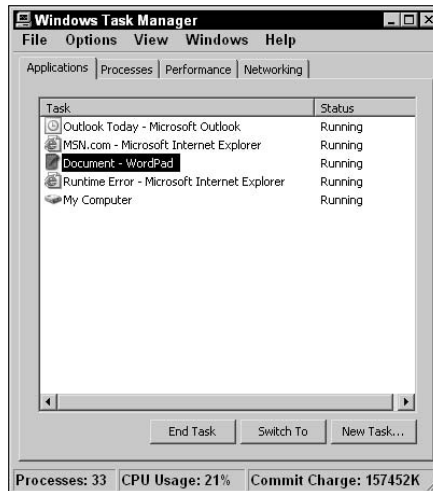


Figure 3-4:
The Applications tab lists running foreground applications and their responsiveness.

Performance

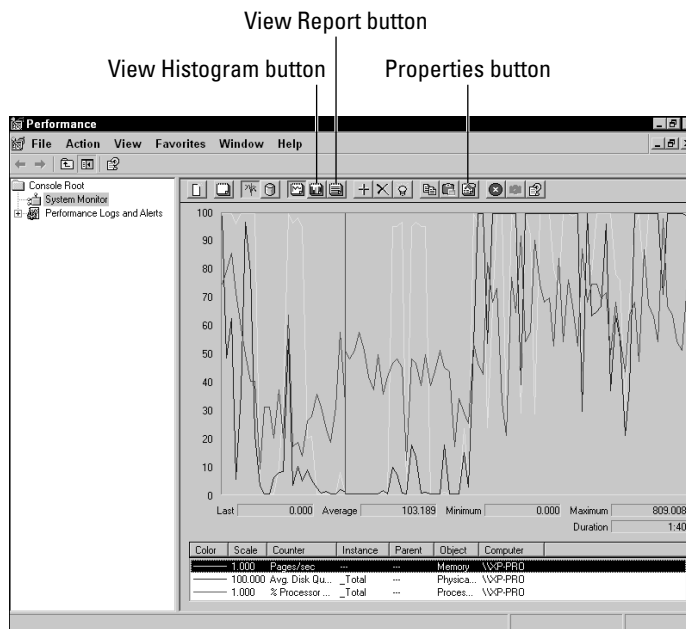
Although Task Manager may be where you first look for solutions to performance problems, it is really a 10,000-foot view of the situation. To get down to ground level, you need a more powerful tool, and that is where the Performance administrative tool comes in, with its two main components, System Monitor and Performance Logs and Alerts.

System Monitor is Task Manager's big brother. It doesn't let you change settings or terminate tasks, but it does let you monitor a whole series of available counters. You can launch Performance by choosing Start⇨Control Panel⇨Performance and Maintenance⇨Administrative Tools⇨Performance, or by running either `perfmon.exe` or `perfmon.msc`. The Performance screen

starts with three default counters loaded into the System Monitor, as shown in Figure 3-5. A counter is a numeric measure of an element of a system component, such as bytes of available memory. The three default counters are:

- ◆ **Pages/sec** from the Memory object to show memory shortage problems
- ◆ **Avg. Disk Queue Length** from the Physical Disk object to show disk access bottlenecks
- ◆ **% Processor Time** from the Processor object to show processor utilization problems

Figure 3-5: Performance Monitor's System Monitor shows graphs for a variety of counters built into Windows.



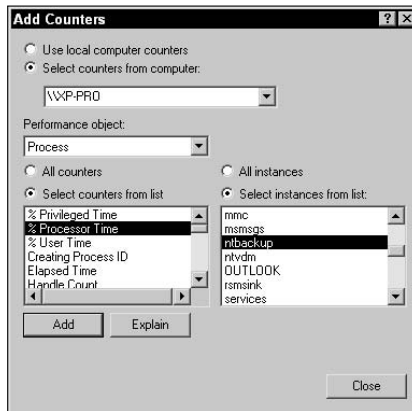
On the exam, you may see references to System Monitor as well as Performance. *System Monitor* is the component of Performance that makes the graphs of current and logged activity.

For each of the counters, you will see a graph with just over 90 readings, which by default are taken once per second. If you want to change this time, you can bring up the Properties for the graph by clicking the Properties button on the toolbar or by pressing Ctrl+Q and changing the value of Sample Automatically Every X Seconds.

If you don't like the line graph look, you can click the View Histogram (think bar graph) or View Report (which gives only numeric data) buttons. These buttons are indicated in Figure 3-5.

Many other counters can be added to the graph. If you click the Add button on the toolbar or press Ctrl+I, you can add counters to your graph, as shown in Figure 3-6.

Figure 3-6:
A wide variety of counters may be added to your graph.



By default, you add counters from *your* computer, but you have the option to add counters for any computer that you have admin rights for. If you prefer to view counters from a remote computer, choose it from the Select counters from the computer drop-down menu to type the computer name into the drop-down menu box. If you choose another computer, your list of Performance objects is updated to include the objects on that computer. Many counters are defaults for the operating system, and any software you install has the option of adding custom counter objects, which is the case for many Microsoft programs. If you look at the list of counters, you will see that they can be categorized into the four critical system resources: processor, memory, disk, and network. Table 3-1 lists some of the related objects for the four critical system resources.

Table 3-1 Related Objects for Critical System Resources

<i>Critical System Resource</i>	<i>Related Object</i>
Processor	Process Processor System
Memory	Memory Paging File

(continued)

Table 3-1 (continued)

<i>Critical System Resource</i>	<i>Related Object</i>
Disk	Cache Physical Disk
Network	Browser IP Network Interface Redirector Server Server Work Queues

Table 3-1 is just a partial list of all of the possible objects that are available to you. Each object has a series of related counters. For instance, the Process object has counters for % Privileged Time, % Processor Time, and % User Time. Each counter may have a series of *instances*. In the case of the Process object and % Processor Time, there are instances for each running process on the system.



For any counters that you do not know, you can select that counter and click the Explain button to get help on that counter.

When you click the Add button to add a counter to your graph, the new counter lines show up immediately. When you have added all of the counters you want to view, click Close.



When you have many lines on your graph, you can click the Highlight button or press Ctrl+H to highlight the line as a white line on your graph for any counter that you have selected on the bottom part of the screen. To delete any counters that you have added, select the counter and click the Delete button on the toolbar or press the Delete key on your keyboard.

So far, you have been taking data from current activity. If the problem is periodic, you want to create logs and alerts to try to catch the problem when it occurs. If you want to create an alert based on your counters, click Performance Logs and Alerts; then right-click Alerts and choose New Alert Settings. You first provide a descriptive name used to identify the alert, and then you can select counters and set the alert-firing threshold for each counter, as shown in Figure 3-7. These values are checked at the Sample interval. When the firing threshold criteria are met, the actions on the Action tab are performed. These include logging the alert in the Event Log, sending a network message to a NetBIOS name, starting a counter log, or running an external program, which might be a script to perform additional actions. This log will be active during the times listed on the Schedule tab.

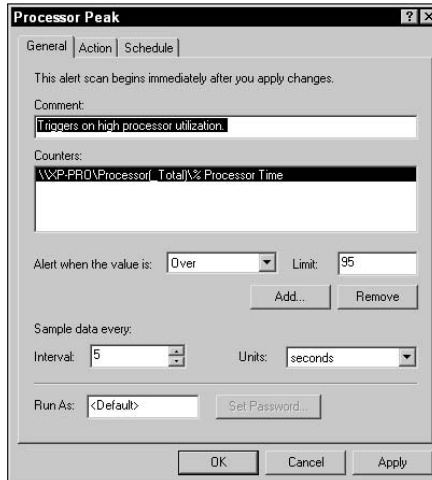


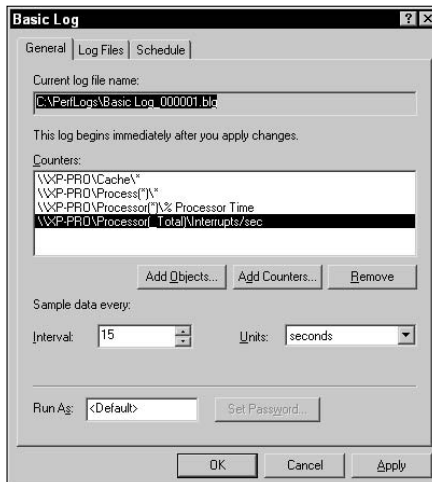
Figure 3-7: Alerts can be used to notify you of problem conditions.

If you want to create a counter log, click Performance Logs and Alerts; then right-click Counter Logs and choose New Log Settings. This opens the window shown in Figure 3-8. Select the counters or objects that you want to log and the interval at which you want to take readings. Then you can set the type of log to create — if you plan to view the data in the Performance Monitor, then you need to use a binary log format. Just like the alerts, you can choose a schedule for this log so that it can run at a specific time. For instance, if you are having problems with a computer regularly between 2 p.m. and 4 p.m., you could schedule the log to run from 1 p.m. to 5 p.m. You can also have the log file stop when the file is full or reaches its configured maximum size. When the logging stops, you can have a program run, which might be used to notify you that the logging has stopped or to copy your new logs to another computer or location.

After a log file has been created, you can view it from the Performance Monitor by clicking the View Log Data button or by pressing Ctrl+L. You are asked for the name of the log file that you want to use, and then you can add counters to the graph normally, but you can choose only from the counters that were included in the log file. The resulting graph charts all readings taken for the counters, not just the default 90 readings.

To change the time range that is displayed on the graph, click the Properties button and the Source tab, from which you can adjust the time range. As you adjust the range, a vertical bar moves through the graph, showing you where that range is. To return to getting graph data from current activity, click the View Current Activity button or press Ctrl+T.

Figure 3-8: Counter logs can be used to record data for historical analysis.



Optimization Best Practices

There are many simple steps that you take to optimize your system for best performance. Some of the components that you can make changes to include virtual memory, hard drives, printers, scanners, system services, running processes, and temporary files. After you have made changes to the components you can use Task Manager and System Monitor to see if your changes have had any effect on the system performance.

Virtual memory

Virtual memory uses both RAM and hard drive space to create a memory pool. The hard drive space that is used is called a paging file, and in Windows XP, the filename is `pagefile.sys`. Access to the paging file is much slower than access to RAM, so the paging file is used for information that is accessed less often. The default paging file size is 1.5 times the amount of RAM in your computer. To improve performance, you should not set this any larger than it really needs to be. You can find out what your maximum size should be by running your system for several days of typical or hard use and then checking your peak usage in Task Manager or System Monitor. Set your paging file size between 1.25 times to 1.5 times your peak paging file usage. To see how to gather your memory usage data using Task Manager or System Monitor, review the “Using Monitoring Tools” section earlier in this chapter; to see how to change the size of your paging file read Book V, Chapter 2; and to get an overview of how virtual memory works, refer to Book V, Chapter 6.

If you find that you're using a lot of virtual memory and accessing the paging file, then you need to either reduce the processes that are using RAM or add more RAM (see Book II, Chapter 3). By reducing processes or adding RAM, you will improve overall system performance since paging file usage will be reduced.

If neither of these steps is an option, then, to improve performance, you should move the page file to a drive other than the drive that holds your Windows installation. You should also choose the fastest drive that you have, so choose 7,200 RPM drives over 5,400 RPM drives, ATA 133 drives over ATA 66, and so on. The chosen drive should not have other highly intensive processes using it.



To avoid having your page file fragmented, you should defragment your drive prior to creating your page file.

To change the location of your paging file, follow these steps:

- 1. Choose Start → Control Panel and open the System applet.**
- 2. Click the Advanced tab and click the Performance Settings button.**
The Performance Options dialog box appears.
- 3. Click the Advanced tab and click the Virtual Memory Change button.**
- 4. From the Virtual Memory dialog box, shown in Figure 3-9, select the drive you want the paging file saved on, choose the size, and then click the Set button.**

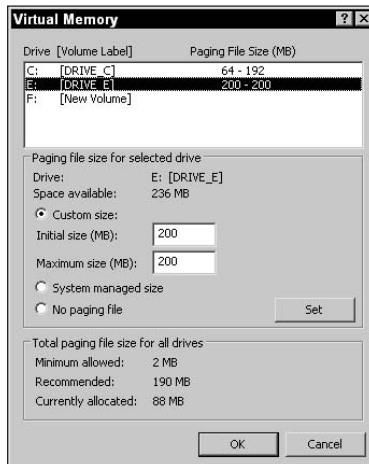


Figure 3-9: Changing the page file settings can improve your overall system performance.

5. To finish off, press the series of OK buttons to close the dialog boxes.

Depending on the changes you made, you may need to reboot your computer.

Depending on the Startup and Recovery setting in your System control panel applet, you may require a paging file on the drive that contains your Windows directory that is at least equal to the amount of RAM installed in your computer.

Hard drive

You can do several things to optimize your hard drives. The first thing is to choose the fastest possible drives for your system. If you can use ATA 133 over ATA 66 or SATA 300 over SATA 150, then you should do so. Keep enough free space on your drives to allow for efficient defragmentation — Microsoft's Disk Defragmenter suggests 15% free space. If you have multiple drives in your system and you are suffering from a disk bottleneck, then you need to move some of your applications from one drive to another to better balance drive utilization.

Printers

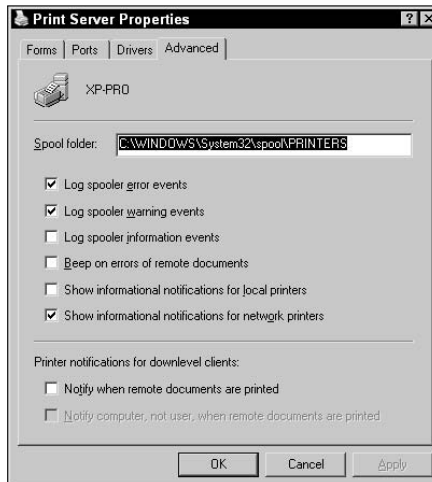
To optimize the printing process, there are a few things that you can do.

- ◆ Make sure you're using current drivers for your printers; upgraded drivers often implement new features and are better optimized for your hardware.
- ◆ Move your spool directory to a different drive. Prior to sending data to the printer, the data is spooled to the computer's hard disk. In order to improve disk access, you should always try to place the spool directory on a fast disk that does not have competition from other applications or services that are running on the computer.

To change the location of the spool directory, choose Start⇒Printers and Faxes, choose File⇒Server Properties, and then click the Advanced tab, as shown in Figure 3-10. From there, you can change the Spool Folder path.

If you are just printing the odd document, then optimizing the printing process isn't an issue. But, if you are using a computer as a print server for an office and interfacing with multiple printers, these steps improve your printing performance.

Figure 3-10:
Changing
the spool
folder
path may
improve
printing per-
formance.



Scanners

Optimizing the scanning process relies mostly on the hardware that is being used. Scanners come with a variety of interfaces, with the oldest being SCSI and parallel. To improve scanning performance, you should choose a scanner that matches the fastest bus architecture (see Book II, Chapter 1) available on your computer — that usually means USB or FireWire. This increases the transfer rate between the scanner and the computer.

The other factor that affects scanner performance is the dots per inch (DPI) level at which you are scanning. If you scan at a lower rate, your scans run faster and the output has a smaller file size — but the quality of the image is lower. The lowest acceptable quality level will be dictated by how you want to use the image. Photographic reproduction requires a higher quality level than does newsprint.

Temporary files

Many applications create temporary files. A temporary file can be thought of as a scratch pad which is a working area for data. An application will use a temporary file to store data that it is working with and then delete the file when the task is completed. For instance, when you open a document in Microsoft Word, Word creates other files in the same directory that start with a ~ character. These other files that Word creates are temporary files, which hold changes to the original document, and automatic recovery information. When you close the document, Word deletes these files, or at least it is supposed to.

It is the job of each application to delete its temporary files when they are no longer needed. Windows uses the environment variables of `%temp%` and `%tmp%` to point applications to the temporary directory. The default location for the temporary directory is in the user's profile, using the path `%USERPROFILE%\Local Settings\Temp`. Having the temporary file directory in the user's profile means that there is a temporary directory for every user on the computer. You can modify each user's temporary directory settings to point to a single location and Scheduled Tasks by running a command like `C:\Windows\system32\cmd.exe /c del /s /q c:\temp` when the computer starts up. Applications are supposed to clean up their files, but after a period of time, you will likely have several files in your temporary directory. The Windows Disk Cleanup utility can delete temporary files, as well as other files not needed on your hard drive.

To get to Disk Cleanup choose Start→My Computer; then right-click on a hard disk →Properties, and click the Disk Cleanup button. A Disk Cleanup dialog box will open, and after scanning your disk the dialog box will change, presenting you a list of items that can be cleaned off your hard disk, such as temporary Internet files, items in the recycle bin, and temporary files. Place a check in the box next to any items you want cleaned, click OK and click Yes in the confirmation dialog box.

Services

Each *service* represents an application that is running on your computer. The more applications that you have running, the fewer system resources you have available for other applications or processes on your computer. To improve overall resource availability and system security, you should disable any services that are not required. You can see a list of all system services through the Services Administrative Tool. To open the Services tool, choose Start→Run and type `services.msc`, or locate it in the Administrative Tools folder, which is found in the Control Panels folder. The Services tool lists:

- ◆ **Name:** The name of service.
- ◆ **Description:** A brief description of what the service does.
- ◆ **Status:** Tells you if the service is currently running.
- ◆ **Startup Type:** Tells you if the service startup is Automatic or Manual or if the service is disabled.
- ◆ **Log On As:** The user credentials used to start the service. Most services start as the Local System, which is the computer's own account.

When you select a service, you will be able to read its description on the left of the window. By reading the description, you may be able to decide whether you need to have that service running. If you want to test to see if

you need a service, turn it off, right-click on the service name and choose Stop. You can tell quickly if you need that service, since something you use will stop working. By turning off services by choosing Stop, the service will restart when you reboot your computer. To have a service remain stopped after a reboot, right-click on a service name and choose Properties; to open the service properties dialog box, choose Manual from the Startup Type drop-down menu, and click OK to close the dialog box.

Prior to randomly turning off services, you can check whether the service is required, by using Internet resources, like Microsoft's Web site, which will provide detailed descriptions of what most services are used for. If you *still* cannot decide whether a service is required, you can disable the service on a test computer and see what happens. Table 3-2 has a summary of the major services that are part of Windows XP.

<i>Name</i>	<i>Description</i>
Alerter	Service to process and deliver administrative alerts.
Computer Browser	Maintains a list of other computers on the network.
Error Reporting Service	Allows error reporting to user and to Microsoft.
Event Log	Logs messages issued by Windows-based programs and components into logs viewed with Event Viewer.
Fast User Switching Compatibility	Provides management for applications in a multiple user environment.
Help and Support	Enables the Help and Support Center to run.
IMAPI CD-Burning COM Service	Manages Windows-based CD recording, rather than using third-party tools.
Indexing Service	Indexes contents and properties of files.
Logical Disk Manager	Detects and monitors new hard disk drives and sends disk volume information to the Logical Disk Manager Administrative Service.
Messenger	Transmits and receives net send and Alerter service messages.
Network Connections	Manages the Network and Dial-Up Connections folder.
Performance Logs and Alerts	Collects performance data from local or remote computers and generates alerts, based on settings found in Performance Logs and Alerts.
Print Spooler	Loads files for deferred printing.
Secondary Logon	Allows starting of applications using an alternate set of user credentials.
Security Center	Monitors system security settings and configuration.
Server	Supports file and print sharing over the network.

(continued)

Table 3-2 (continued)

<i>Name</i>	<i>Description</i>
System Restore Service	Performs system restore functions based on System Restore Points.
Task Scheduler	Enables configuration and scheduling of automated tasks.
Terminal Services	Base multiuser components that are used by Remote Desktop, Fast User Switching, Remote Assistance, and Terminal Server.
Themes	Manages XP general desktop themes.
Volume Shadow Copy	Allows Volume Shadow Copies that are used for backup and other purposes.
Windows Firewall/ Internet Connection Sharing (ICS)	Provides firewall and Internet gateway services, such as name resolution, network address translation, and intrusion prevention services.
Windows Installer	Base Windows component that allows for the installation of .msi files.
Wireless Zero Configuration	Provides a standard configuration interface for 802.11 adapters.
Workstation	Creates and maintains client network connections to remote computers.

There are properties that can be set and managed for services. If you want to change the properties of a service, you must open the Services Administrative Tool, as previously mentioned; then locate and select the service you want to modify or view settings for and then right-click Properties. This opens the service properties dialog box, which has three tabs: General, Log On, Recovery, and Dependencies General tab.

This tab lets you see basic information for the service, and, most of the time, this is the only tab you will need to use. This will show you the display name that is used in the Services tool, as well as the path to the executable that is run to start the service. There is a startup type that can be set to Automatic, Manual, or Disabled. The Automatic settings will start the service on computer reboot; Manual will allow it to be started using the `net start` command or the Services tool; Disabled will not allow the service to be run at all.

There are also service control buttons to start, stop, pause, and resume a service. Start and Stop are self-explanatory and are supported by all, but some services support pause and resume. Pause will typically prevent new requests from being processed by the service, but will allow existing requests to be processed. Resume will restore a paused service to normal operation. One case in which you might use this is with the Server service, which allows people to access files on your computer from across the network. Stopping the service will disconnect all users, while pausing the

service will prevent new users from accessing files on your computer, but people who already have that connection open will be able to finish their work. Once all users have completed their work and have disconnected from your computer, you can complete what you were intending to do, which might be to stop the service, make a configuration change, and restart the service.

Some services accept start-up parameters in the same way that you can pass parameters to other applications. When a service is stopped, you can add or change the Start parameters, and then start the service.

Log On tab

All services are programs, and all programs on a Windows NT-based computer will run using security settings of a user account. Most services run using the operating system's account, also called Local System; however, you may want a service to run using a different account so that you can restrict what the service can do or so that the service can interact with other computers on the network.

In addition to account settings, the Log On tab allows you to specify which hardware profile this service should run for. This allows you to use hardware profiles as tools to control which services will be loaded during any reboot. Hardware profiles are covered in Book VI, Chapter 1.

Recovery tab

The Recovery tab allows you to deal with what to do when the service stops unexpectedly. Three drop-down menus enable you to specify an action for the First failure, Second failure, and Subsequent failures. For each of these menus, you can choose Take No Action, Restart the Service, Run a Program, or Restart the Computer. The default is to take no action.

If you have a service that fails for an unknown reason, you can set the service to restart on failure while you are trying to find the root cause, which may be enough to continue to provide the functionality required.

If you specify to run a program, then you are able to use the bottom of the Properties dialog box to specify what program or batch file you want to run and the parameters to pass to the program. The program or batch file may be used to automatically fix a known problem or to send an e-mail or alert to the administrator of the computer.

If you choose to restart the computer, you can click on the Restart Computer Options button at the bottom of the dialog box to set the delay for the reboot, and a message to send to people connected to the computer.

Finally, the Recovery tab enables you to specify the number of days at which you will reset the failure counter and the time to wait before restarting a service.

Dependencies tab

Some services require that other services are running, prior to starting up. If you try to start a service that depends on other services that are not currently running, the required services will also be started. The Dependencies tab allows you to see both services that the selected service requires, as well as what other services need the selected service.



One service that I often disable on systems that have limited resources is the Themes service. This service is responsible for drawing windows with rounded corners and most of the Windows XP graphical enhancements. All of these items take processor time and memory to render for the user. By disabling the Theme service, these enhancements are automatically disabled. There may be some cases in which a user might feel that this service is required, and the same can be said of any service. For instance, an anti-spyware service running in real-time mode, checking each file that is accessed on your hard drive, is a good service to have running on a computer that is connected to the Internet and is used for Web browsing, but this same service may be of limited use on an isolated computer with no Internet access.

Startup

When your computer starts up, it loads all of its services as well as any applications that are referenced in the Run registry keys and the Startup group in your start menu. The Run registry keys include:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Run  
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\RunOnce  
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\RunOnceEx  
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Run  
HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\RunOnce
```

These items represent applications that are running on your computer. Many of these items, as well as some services, are responsible for the ever-growing string of icons next to your clock in the system tray. Each of these is a running application or process that takes system resources away from other applications or processes running on your computer.

Many of the icons in your system tray have preferences or options that allow you to stop the background process from running. If you disable this service, it may mean that some applications take longer to load. For instance, Sun's Java or Apple's QuickTime use the startup registry key to place one of their

startup applications in the system tray, and they use the application to pre-launch their main application environments, so when you launch a QuickTime movie, QuickTime is already running and only needs to open the media file. If this process were not running, then you would have to load QuickTime into memory before running the media file. On systems with limited resources, you will want to remove many of these pre-loaded components to free the critical resources they use. Even though each icon represents a small amount of resource, they all make up the straws that are on the camel's back, and you never know where the breaking point is.

Getting an A+

This chapter reviews optimization of the Windows environment. Major points covered in this chapter include:

- ◆ Virtual memory and applications that access the paging file reduce performance.
- ◆ Performance is increased by using the computer's faster interface available for a device, such as using FireWire rather than serial.
- ◆ To improve performance, you should reduce the number of processes and services that are running on the system.
- ◆ Printers can be optimized by changing the location of the spool directory.

Prep Test

- 1** What tools would you use to monitor disk I/O? (Choose all that apply.)
 - A replmon
 - B perfmon
 - C defrag
 - D taskman

- 2** Virtual memory contains which item?
 - A A special section of memory used for caching data
 - B An area of a hard drive for caching data
 - C An expanded memory PCI expansion card
 - D An extended memory PCI expansion card

- 3** What tool manages your virtual memory settings?
 - A Memory control panel
 - B setver command
 - C System control panel
 - D mem command

- 4** What two steps can you take to improve hard disk access? (Choose two.)
 - A Double the number of pins on your drive that are carrying data.
 - B Defragment your drive.
 - C Add a speed doubler to the drive bus.
 - D Change to a faster bus architecture.

- 5** Which of the following will *not* help optimize printer performance?
 - A Using a newer driver
 - B Moving the spool directory to a faster drive
 - C Reducing the number of other applications running on the print server
 - D Allocating more RAM to the printer cache

- 6** Overall system performance can be improved by all of the following except for which item?
 - A Defragment drives.
 - B Add more memory.
 - C Reduce the number of processes.
 - D Increase the page file size.

Answers

- 1 B, D.** Both Task Manager and Performance Monitor can be used to monitor disk input and output. Performance Monitor has more detailed results, while Task Manager provides only a few basic counters. *See "Using Monitoring Tools."*
- 2 B.** Virtual memory is created by using a page file, which resides on a hard drive. Disk caching uses memory to cache data from the hard drive. *Review "Virtual memory."*
- 3 C.** Virtual memory settings are managed through the System control panel. *Check out "Virtual memory."*
- 4 B, D.** Two things that can be done to speed up disk access are defragmenting your hard drive and changing to disks that have faster bus architectures. *Peruse "Hard disk."*
- 5 D.** There is no specific area of RAM that is used for caching printer data. *Take a look at "Printers."*
- 6 D.** Increasing the size of the page file will not provide a performance benefit. *Peek at "Optimization Best Practices."*

Chapter 4: Using Windows-Based Troubleshooting Utilities

Exam Objectives

- ✓ Managing registry data by using `regedit.exe` and `regedt32.exe`
- ✓ Working with system and disk management tools

This chapter is all about Windows-based troubleshooting utilities. These utilities fall into categories such as disk maintenance, configuration, diagnostic, and the ever-popular miscellaneous category. One major utility topic covered in this chapter is Registry-editing utilities. I examine this component thoroughly because this topic is often neglected. The content that you examine for Registry editing is more detailed than what you need for the exam; when you finish this book, you'll have a great grasp on the capabilities of the Registry editing tools.

As a CompTIA A+ Certified Professional, you will be called on from time to time to use the tools that are discussed in this chapter. So, in addition to knowing this content for the exam, you need to understand these tools to do your job as a support professional.

Understanding Operating System Utilities for Disk Maintenance

In this section, you examine different utilities that you can use to perform disk maintenance on your system. These utilities include `fdisk` and Disk Management, `defrag`, and `chkdsk`. Each of these utilities serves a specific purpose for your operating system.

fdisk.exe and Disk Management

Before you can store any data on your disk, you need to prepare the disk for use. You can use the `fdisk.exe` utility to prepare your disk to hold data if you are using an MS-DOS boot disk, or you can use Disk Management (`diskmgmt.msc`) if you are in Windows.

Data is stored on a disk in a partition, and that partition is assigned a drive letter. Because of a decision made when early computer disks were created, you are allowed to create only four partitions on a disk, and a table was created in the disk's BIOS to hold the partitioning information. The partition table has only four entries in it. The solution to the four-partition limit was to create a special partition, which is always stored in the last partition. This special partition is called an *extended partition*, and it is designed to extend the number of drive letters you can associate with a drive. Remember that this is four, the same as the number of partitions. The extended partition does not have a drive letter associated with it, but rather allows you to create any number of logical drives, which do have drive letters.

The `fdisk.exe` utility dates back to the earliest versions of MS-DOS, in which it was used to create partitions on your disk. The `fdisk.exe` utility was modified with the release of Windows 95 OSR2 to allow it to support FAT32 partitions. The `fdisk.exe` utility supports many command-line switches as well as an interactive interface. Most people are familiar with the interactive interface, so I cover that first and the command-line switches later in this section.

Windows 2000 and Windows XP do not use `fdisk.exe`; they use a program called Disk Management, which is a *snap-in*, or a tool component for the Microsoft Management Console. Disk Management is fully covered in Book II, Chapter 5.

Knowing your way around `fdisk.exe` is still important because you may have to resort to using `fdisk.exe` to prepare a driver for disk imaging or installing Windows XP in certain configurations. After booting your computer from an MS-DOS or Windows 9x boot disk — you will be sitting at a command prompt — type `fdisk.exe` to launch the disk partitioning tool. If you are using `fdisk.exe` from either Windows 95 OSR2 or Windows 98, then you will be prompted on whether to enable large disk support. If you enable this support, then any new partitions that you create in excess of 512MB will be formatted as FAT32, leaving them unreadable by operating systems that can't read FAT32 partitions, like Windows NT and most versions of Windows 95.

After deciding whether you would like to support FAT32, you'll be presented with the normal list of options for the `fdisk.exe` utility. These options include:

- ◆ **Create DOS partition or Logical DOS Drive.**
- ◆ **Set the active partition — your drive will not be bootable without one.**
- ◆ **Delete partition or Logical DOS Drive.**
- ◆ **Display partition information.**

Each drive can have only one primary partition. You can also create one extended partition on each drive, which can contain multiple logical drives. Primary partitions are represented on your computer with an associated drive letter, while extended partitions do not have drive letters, but rather contain logical drives that have associated drive letters. Each of the primary partitions or logical drives that appear on your computer is still limited to the size that is specified by the file system. This means that FAT16 drives are limited to 2GB, and FAT32 drives are limited to 2TB.



You can boot only from an active partition, so if you don't set up an active partition on a new drive, then you won't be able to boot from that drive. If you select No in response to the prompt to create a partition of the maximum size and make it active, then you will have to create the active partition yourself.

If you use Disk Management, you can create up to four primary partitions on each drive. If you want to create an extended partition, then you can create only three primary partitions. You also have the option of updating your disk to a *Dynamic Disk*, which eliminates the traditional limits to partition table entries by maintaining its own partitioning information in a separate location on the disk.

Dynamic Disks do not have as many recovery tools as traditional Basic Disks and are only compatible with Windows 2000 and newer Microsoft operating systems. Due to this recoverability aspect, you should avoid Dynamic Disks unless they are necessary for your disk configurations.

Figure 4-1 shows the partition table of a typical 6GB hard drive. The hard drive contains a 2GB partition and a 4GB partition. The 4GB partition will likely contain two or more logical drives. If the extended partition contains logical drives, `fdisk.exe` asks whether you would like to see them.

Figure 4-1: Primary partitions and logical drives are associated with drive letters.

```

Display Partition Information

Current fixed disk drive: 1

Partition Status Type Volume Label Mbytes System Usage
C: 1 A PRI DOS 2055 UNKNOWN 33%
2 EXT DOS 4007 UNKNOWN 67%

Total disk space is 6142 Mbytes (1 Mbyte = 1048576 bytes)

The Extended DOS Partition contains Logical DOS Drives.
Do you want to display the logical drive information (Y/N).....? [Y]

Press Esc to return to FDISK Options
    
```

In the Disk Management utility, the entire contents of the disk partitioning scheme are visible in the display, and you are given more control over the drive letters that are assigned.

The `fdisk.exe` utility supports several switches, but only a few are actually documented. One of the more useful and heavily documented of the switches is `/status`. This switch returns a one-screen listing of your drive's partition tables and then exits the `fdisk.exe` program.

The `/mbr` switch is one of the undocumented `fdisk.exe` switches. This command rewrites the *Master Boot Record (MBR)* for your hard drive. The MBR is read by your computer's *POST (Power-On Self-Test)* process when it is looking for a bootable drive. If your drive should be bootable but is not being detected as bootable by your computer, you can try rewriting the MBR. This may give you a bootable drive. This process may also help to remove some viruses that store themselves in the MBR, but a good antivirus program provides better protection and removal.



Take care when using the `/mbr` switch with `fdisk.exe` because a number of programs may have problems if the MBR is rewritten. Rewriting the MBR can also cause problems if your computer is currently configured to dual boot with another operating system that has modified the boot sector, such as Windows 2000 or Windows XP.

The last of the undocumented options allows you to partition your hard drive without loading the interactive version of `fdisk.exe`. The correct syntax for this command is:

```
FDISK.EXE <drive_number>/<Pri|Ext|Log>:<Size in MB>
```

If you want to partition a new drive into a primary partition of 2,048MB, an extended partition containing a logical drive of 2,048MB, you can use the following commands:

```
FDISK.EXE 1/Pri:2048
FDISK.EXE 1/Ext:2048
FDISK.EXE 1/Log:2048
```

Rather than using the command-line switches, you can run `fdisk.exe` interactively, which is required if you need to delete or change the partition table on your hard drive.

defrag.exe and dfrq.msc

In this section, I discuss only the defragmentation options available to you in Windows. I exclude MS-DOS-related defragmentation because you will only be tested on the concepts of defragmentation or on the methods used to

defragment disks within Windows. The MS-DOS defragmentation program has the same name as the Windows command-line version, `defrag.exe`, and the MS-DOS version is self explanatory if you every find yourself in a position to run it.

As you write data to your hard drive, it is always written into the largest area of free space that is available, but when you delete files, areas of free space become broken up and scattered around your drive. Sometimes the largest area of free space isn't large enough to hold the data that is being written. When this happens, the file is broken up into pieces to the areas of free space.

When files are broken into pieces, they are *fragmented*. You can use `defrag.exe` or `dfrg.msc` to "repair" the disk fragmentation. `dfrg.msc` is the graphical defragmentation tool for Windows, while `defrag.exe` is a command-line front end for the same tool. Both of these tools are shown in Figure 4-2. If you want to create a batch file or script to defragment your hard drive, you want to look at `defrag.exe`.

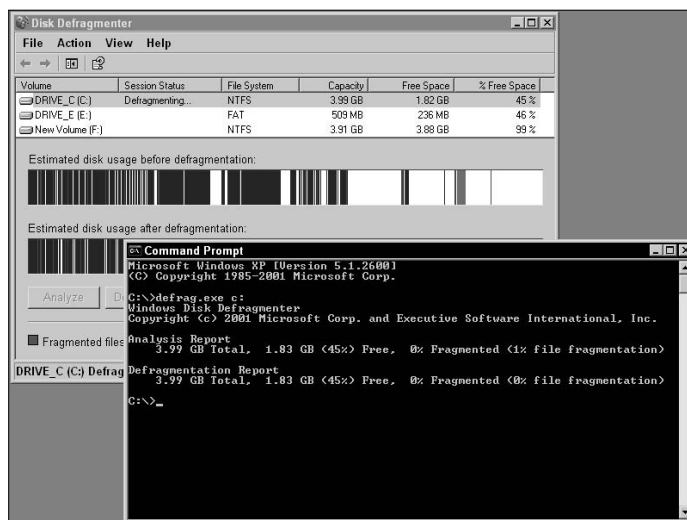


Figure 4-2: You have two interfaces to choose from in order to defragment your drives.

Defragmentation requires free space to do its job because it copies the files to new locations before removing them from the old locations. The defragmentation utility does not touch files that are open. In addition to making all files contiguous, defragmentation attempts to consolidate most free space into a single unit, reducing how often new files are fragmented.



While the defragmentation process takes place, all files on your drive are read, and the system is checked for errors, in the same manner as `chkdsk.exe` (discussed in the next section).

Because the disk fragmentation utility rearranges the data on your drive and can therefore improve performance, it should be run periodically to ensure maximum performance of your computer. Just what *periodically* translates into depends on your system and how you use it. If your hard drive is extremely full (containing perhaps only a few gigabytes of free space), then you should run `defrag.exe` as often as every few days. If you have a large amount of free space on your drive, files will take longer to become fragmented on your disk, and you may need to defragment your drive only once every month or two.

To open the graphical disk defragmenter, choose Start⇨All Programs⇨Accessories⇨System Tools⇨Disk Defragmenter. When you select a drive and click the analyze button, the defragmenter recommends whether or not you need to defragment your drive. But if you are noticing a performance degradation, your drive should be defragmented by selecting a drive and clicking the Defragment button. There is no harm in defragmenting your drive, even if the computer says you don't need to; you may still increase the computer's performance.



You can use Scheduled Tasks to regularly run disk defragmentation. I cover scheduling tasks in the “Miscellaneous Utilities” section of this chapter.

The important thing to remember is that free space and disk performance should be monitored in order to achieve an optimal level of performance.

chkdsk.exe (Check Disk)

The `chkdsk.exe` utility can identify potential problems with either your file system's allocation table or index, or the physical hard drive. `chkdsk.exe` only comes in a command-line interface, but it has several options, of which the most important are:

- ◆ /F: Fixes errors found on the disk, rather than just identifying them.
- ◆ /R: Relocates bad sectors and recovers readable information.
- ◆ /X: Forces a dismount of the volume, which closes all open files.

If you suspect that you have a disk problem, you should run a scan of your disk; but because fixing errors requires exclusive access to the disk, you want to run `chkdsk.exe` in scan-only mode, which is the default. If errors are detected, then you would want to specify /F and /X to correct any errors that are identified.

The system volume (C:) cannot be dismounted, so `chkdsk.exe` offers to schedule the scan during the next reboot of the system. On a system with a large C: drive, and depending on the number of errors, this scan can be very time-consuming. The results of the scan can be seen Figure 4-3.

Figure 4-3:
Standard scans will check the structure of the file tables.

```

C:\>chkdsk.exe C:
The type of the file system is NTFS.
Volume label is DRIVE_C.

WARNING! F parameter not specified.
Running CHKDSK in read-only mode.

CHKDSK is verifying files (stage 1 of 3)...
File verification completed.
CHKDSK is verifying indexes (stage 2 of 3)...
Index verification completed.
CHKDSK is verifying security descriptors (stage 3 of 3)...
Security descriptor verification completed.
CHKDSK is verifying Usn Journal...
Usn Journal verification completed.
Correcting errors in the Volume Bitmap.
Windows found problems with the file system.
Run CHKDSK with the /F (Fix) option to correct these.

4184980 KB total disk space.
2213224 KB in 14678 files.
3772 KB in 952 indexes.
0 KB in bad sectors.
45396 KB in use by the system.
22976 KB occupied by the log file.
1917588 KB available on disk.

4096 bytes in each allocation unit.
1846225 total allocation units on disk.
479327 allocation units available on disk.

C:\>
    
```

In addition to issues with directory index entries and *MFT (Master File Table)* problems, `chkdsk.exe` looks for cross-linked files and orphaned directory entries or lost file fragments. The file allocation table contains a pointer to the starting cluster for each file. As the file is read, a pointer at the end of each cluster identifies where the next cluster of the file can be found, with the last cluster of the file not containing a further pointer. From time to time, the pointer at the end of a cluster becomes corrupted and starts pointing to the wrong location for the next bit of data. If the new location already contains file data (which contains a pointer to the next bit of data for a second file), you end up with two files using the same set of clusters.

When two files point to the same cluster (and clusters that follow), the situation is known as a *cross-linked file*. Depending on your outlook, you may choose to make copies of these files or to ignore them altogether. If you are a pessimist, then you probably believe that both of these files are corrupt and probably garbage. If you are an optimist, then you probably believe that one or both of these files are still readable and usable. There is no harm in having the check disk create the files, as your worst case scenario will be that they files are of no use.

File fragments or lost chains, on the other hand, tend to be the leftover portions of cross-linked files (the portions of a file that would have belonged to a file had its clusters not been corrupted or lost). You have the option of either deleting these file fragments or converting them to files. If you choose to convert them to files, then each chain is saved to the root directory with a name following the format `filennnn.chk` with an incremental number in place of the `nnnn`.



Although you have the option to convert lost file fragments to files, I have yet to see useful data recovered from file fragments, so you might as well delete them and free the space.

`chkdsk.exe` serves an important purpose on your system by ensuring data integrity of the files on your disk, and it should be scheduled to run periodically on important volumes because some disk corruption can be prevented if identified early.


Configuration Utilities

This section discusses many of the configuration utilities that ship with the Windows operating system. These utilities provide easy ways to modify configuration files and the Registry.

System Configuration Utility

The System Configuration Utility (see Figure 4-4) is one of the newer utilities for maintaining settings on your Windows computer. To launch the System Configuration Utility, run `msconfig.exe`. This utility can change many aspects of your system startup. In order to use this utility, you need to be a member of the computer's Administrators group.

Figure 4-4: The initial window for the System Configuration Utility allows you to modify the basic system files.



The six tabs in this utility allow you to launch the System Restore utility, expand files, perform selective boots, modify the loaded drivers in `system.ini` and `win.ini`, change the boot options in `boot.ini`, and change the services and applications that are loaded at boot time.

This application is capable of doing the following:

- ◆ **Creating custom startup configurations (for troubleshooting)**
- ◆ **Performing a selective startup, which only processes some of the system files**
- ◆ **Editing `boot.ini`, `system.ini`, and `win.ini` files**
- ◆ **Expanding operating system files from the OS CD**
- ◆ **Disabling services or programs that have been added to the Run key of the Registry**

All the changes in the preceding bulleted list can be accomplished from within a simple Windows user interface, making this an easy configuration tool to work with.

If you think a computer problem is related to one of the components that are loaded at startup, then this is the tool you want to use. If you can boot into Safe Mode but not perform a normal boot, you can use this tool to select which components you want to load on reboot. By disabling components, you should be able to narrow down which component is causing the problem, by performing multiple boots with different combinations of the items selected.

regedit.exe

If you have ever talked to anybody about the system Registry, then you have probably been given a stern warning that this is not a place for the timid. The reason for this warning is that the Registry stores system-wide configuration information for almost all services on your computer. If you experience major problems with the Registry, you will likely have major problems with your entire computer.

In Windows 3.1, the Registry held information about file associations. With the creation of Windows 9x, Windows NT, and more recent Windows operating systems, the Registry now stores key information about all system services.

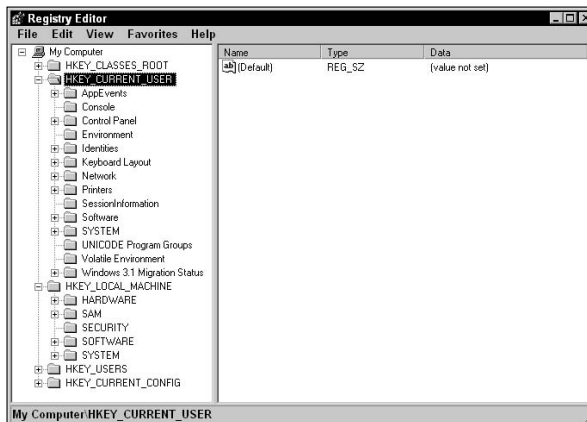
The Windows Registry has five major sub-trees, as shown in Figure 4-5:

- ◆ `HKEY_CLASSES_ROOT` contains information about all the file associations and registered file types that exist in your computer. This information is actually a copy of the information stored in the `HKEY_Local_Machine\software\classes` key.
- ◆ `HKEY_CURRENT_USER` contains a subset of the information that is found in `HKEY_USERS` key, but only contains the information related to the current user. This key contains configuration information for the currently

logged-on user. This information includes items like Control Panel settings (such as mouse acceleration and screensaver preferences) and user-related software settings for installed applications.

- ◆ HKEY_LOCAL_MACHINE contains information about your computer. This information includes configuration information for hardware components and system settings for all software on the computer.
- ◆ HKEY_USERS contains information for the currently logged-on user, as well as information for the default user and all other users who are logged onto the computer.
- ◆ HKEY_CURRENT_CONFIG contains settings that are dynamically rebuilt on every boot of the computer, which is mainly composed of plug-and-play information.

Figure 4-5:
Settings in
the Registry
cover most
aspects of
the system.




Of these sub-trees, you really need to worry about only two: HKEY_CURRENT_USER and HKEY_LOCAL_MACHINE. The information in the other sub-trees is either duplicated within these two sub-trees or the data cannot be modified.

All the entries in the Registry are stored in a hierarchical tree. This tree contains several subfolders that are called *keys*. If you want to create new keys within the Registry, you can do so by using the Edit→New→Key command. Every key within the Registry is capable of storing values. The Registry Editor contains five different types of values (as shown in Figure 4-6):

- ◆ String
- ◆ Binary
- ◆ DWORD

- ◆ Multi-String
- ◆ Expandable String

Figure 4-6:
There are five basic data types for the Windows Registry.



Some hardware components or software components that you can install on your computer may create entries inside both `HKEY_LOCAL_MACHINE` and `HKEY_CURRENT_USER`. This makes sense if you think about applications like Microsoft Office. Microsoft Office creates entries in `HKEY_LOCAL_MACHINE` that relate to the location of Office components on the particular computer, such as the spell checker that can be found within `C:\Program Files\Microsoft Office\Tools`. At the same time, Microsoft Office installs entries in `HKEY_CURRENT_USER` that relate to the user configuration on the system. User-related settings are items such as what the default Save As file type will be when you use Microsoft Word.



The Registry is an important part of your computer, and great care should be taken not to corrupt this database. Before you start doing any procedure that involves the Registry, make sure you have a clean backup of the Registry. You can create this backup in a variety of ways, such as by exporting the Registry, by creating a System State backup, or by physically copying the files that make up the Registry.

To copy the files that make up the Registry, you need to know where the files actually are. `HKEY_CURRENT_USER`, the user portion of the Registry, can be found in your user profile directory. It will be named `NTuser.dat`. `HKEY_LOCAL_MACHINE` is actually composed of several files that are found in the `Winnt` or `Windows` directory. These files should be backed up and restored by using the Windows backup program. The Windows 2000 backup utility is capable of creating an Emergency Repair Disk, which can restore your Registry; with Windows XP, you can perform a System State backup.

Rather than backing up the entire Registry, you can be more selective and back up and restore individual keys within the Registry by using the registry editor.

To back up individual keys, follow these steps:

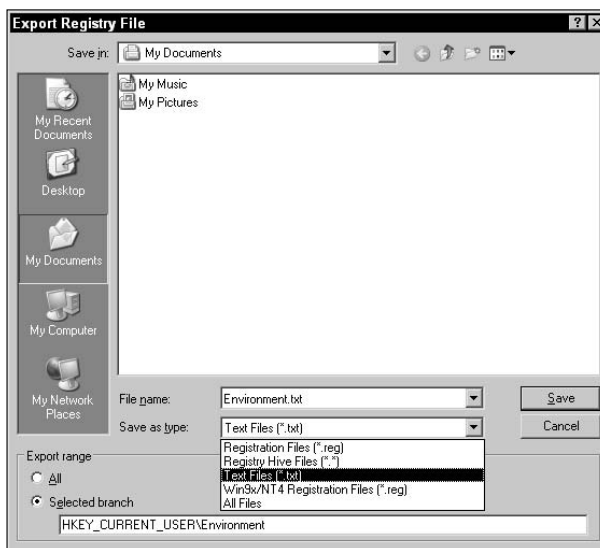
1. Choose **Start**→**Run** and type `regedit` in the Run dialog box that appears.
2. Click the key that you want to back up to highlight it.
3. Choose **File**→**Export**.

Doing so displays the Export Registry File dialog box, as shown in Figure 4-7. This dialog box looks similar to a standard Save As dialog box — with the addition of the Export Range panel at the bottom. The Export Range panel contains your selected branch within the Registry by default.

4. Type a name for the file and choose a location for `regedit` to save it.
5. Click **Save**.

The file will be given the default `.reg` extension.

Figure 4-7: Exporting sections of the Registry is a great way to back up settings in the portion of the Registry that you are editing.



You can use Windows 2000's `regedt32.exe` to save Registry keys to be restored at a later date, but these files are in a binary format and cannot be edited with a text editor.

Files that end in the `.reg` extension are Registry export files that you can edit with a text editor (as shown in Figure 4-8). You should note that the `.reg` files share the same basic structure as an `.ini` file, which you were introduced to in the `system.ini` section of Book V, Chapter 6. To import a

.reg file into the Registry just double-click on the file. To help prevent accidental importing of Registry files, Windows 2000 and Windows XP have added a safety feature: Windows prompts you to confirm that you would like to import the settings into your Registry, as shown in Figure 4-9. After you import the settings, a dialog box confirms that a file has been successfully merged into your Registry.



Figure 4-8: All Registry export files share a structure that resembles .ini files.

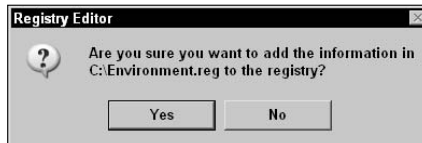


Figure 4-9: The confirmation for Windows 2000 and Windows XP Registry imports helps to limit accidental imports.



On Windows computers, .reg files are automatically imported into your Registry if you double-click them. Rather than using the .reg extension on your files, you may choose to use .txt files. Files with the .txt extension can also be imported into your Registry by typing `regedit.exe filename.txt` at the command line.

Periodically, `.reg` files become corrupted from being edited with certain text editors, such as Notepad. This corruption occurs because some text editors add extended (invisible) characters into the text file that are not compatible with `regedit.exe`. The only editor that does not exhibit this behavior is `edit.com`. If you have a file that has been corrupted, you can fix it by opening the file through `edit.com`, resaving the file, and closing. Even if the file is corrupt, you still receive a dialog stating the information has been successfully entered into your Registry, as shown in Figure 4-10.

Figure 4-10:
A successful Registry import.



The `regedit.exe` program is capable of editing the Registry for Windows 9x, Windows NT, Windows 2000, and Windows XP computers. One of the few differences between the different Registry editing programs is that the Windows 2000 version has added a Favorites menu, which allows you to quickly return to common Registry keys for editing.

regedt32.exe

The `regedt32.exe` program is the Registry Editor for Windows NT and Windows 2000 (but not Windows XP). Most of the user interface for `regedt32.exe` is still based on the Windows NT 3.51 user interface. That means that many of the windows have a Windows 3.0 look and feel to them (see Figure 4-11). The `regedt32.exe` program supports many features that are not available in `regedit.exe`. These features include:

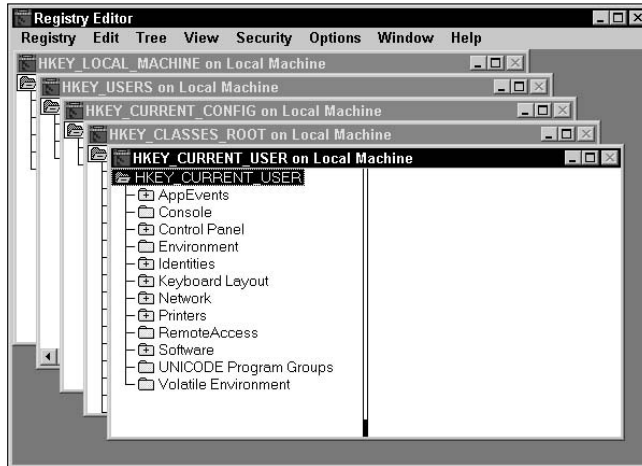
- ◆ **Viewing and setting permissions on Registry keys**
- ◆ **Loading and unloading Registry *hives* (which are Registry files on your hard drive)**



With Windows XP, all features of `regedt32` have been incorporated into `regedit.exe`, and if you run `regedt32`, `regedit.exe` will launch instead.

When working with Windows NT 4.0, `regedt32.exe` allows you to connect to remote Registries, while `regedit.exe` does not. When you work with Windows XP, you can use either program to connect to a remote Registry. Connecting to a remote Registry allows you to change the Registry settings on a remote computer without having to leave your own computer. To connect to a remote Registry from across the network, choose Registry → Select Computer in `regedt32.exe` or Registry → Connect Network Registry in `regedit.exe`.

Figure 4-11:
The
`regedt32`
.exe
program
sports the
Windows
3.0 user
interface.



One big advantage of `regedt32.exe` is that you can load Registry hives. This ability becomes useful when you want to modify the entries for the default `ntuser.dat`. The default `ntuser.dat` is used as a template to create new user profiles for users who do not yet have a profile. You can find the `ntuser.dat` file for this default template account in `C:\Documents and Settings\Default User\NTUSER.DAT`, and you can load this sub-tree or hive into `regedt32.exe` and edit the settings for this default template.

This default user template should not be confused with the default user that is in use prior to a user logging on to your computer. The Registry settings for the pre-logon default user are actually found in the Registry under `users\DEFAULT`.

The other large advantage of `regedt32.exe` is that you can manage the security on Registry keys. In order to view or modify the security settings on a Registry key, first select the key and then choose `Security`→`Permissions`. You can work with permissions on Registry keys only — not on Registry values.

To allow for command-line editing of the Registry, Microsoft introduced `reg.exe` as a standard component of Windows XP. This allows for loading and unloading of Registry hives as well as editing registry keys and values.



Lab 4-1 will have you export your current Desktop settings into a file, and create a registry file that is used to change wallpaper settings.

Lab 4-2 will modify the Desktop settings for the user environment that is in effect when no users are logged on. This is the wallpaper and color scheme that you see behind the logon screen.

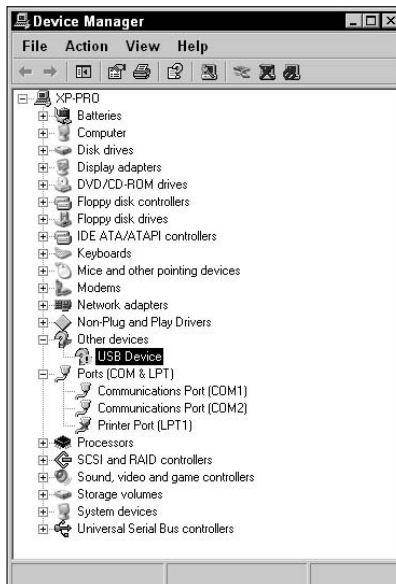
Lab 4-1 and Lab 4-2 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Device Manager

Device Manager is covered in more depth in Book VI, Chapter 1. Device Manager has a primary job of, appropriately enough, managing devices. You can access Device Manager by clicking the Device Manager button on the Hardware tab of the System Properties window. Device Manager lets you remove devices from your system, configure drivers, configure hardware resources, and manage hardware profile settings.

In the Device Manager, you can easily pick out the devices that have hardware problems — they have a red X on the icon in the device tree, as shown in Figure 4-12. The red X identifies devices that are in conflict or are disabled, while a yellow exclamation mark identifies devices that are not properly configured. You can use Device Manager to identify problems with your hardware, configure the drivers for the devices, and configure hardware resources, such as IRQs and I/O addresses.

Figure 4-12: If you are looking for malfunctioning devices, look for the red X and the yellow exclamation mark.



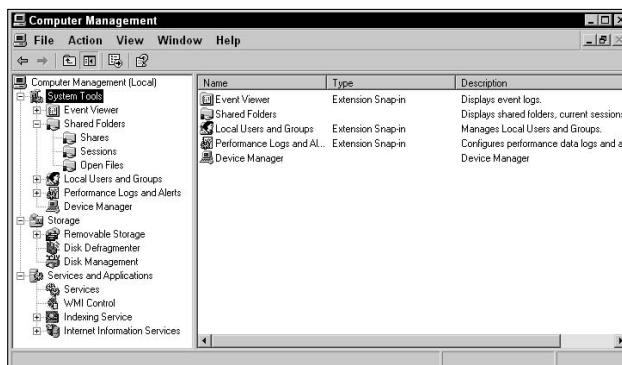
Computer Management

Windows 2000 and Windows XP include a utility called Computer Management, shown in Figure 4-13. Computer Management is actually more of a user interface to a number of other utilities than an actual utility.

Computer Management is a *Microsoft Management Console (MMC)* settings file that contains the following utilities or snap-ins:

- ◆ **Event Viewer:** Allows you to view the contents of event logs files.
- ◆ **Shared Folders:** Shows status of shared or published folders on the computer.
- ◆ **Local Users and Groups:** Manages the local account database.
- ◆ **Performance Logs and Alerts:** Records system status and health based on measurable counters and generates alerts when the counters exceed threshold values.
- ◆ **Device Manager:** Displays hardware status and configuration information in a format that can be edited, unlike the System Information Tool (see the section, “The System Information Tool,” later in this chapter).
- ◆ **Disk Management:** Manages the disk partitions. The Windows 2000 or Windows XP version of `fdisk.exe`.
- ◆ **Disk Defragmenter:** Optimizes disk performance by defragmenting files.
- ◆ **Removable Storage Management:** Manages and tracks removable media on your computer, such as CD-ROMs and magnetic tapes.
- ◆ **Services:** Manages system services (such as the Server and Workstation services) on the local computer.
- ◆ **Other administration utilities:** Other tools are included with this tool — for example, Telephony, WMI Control, Indexing Service, DNS, Internet Information Services, and Routing and Remote Access.

Figure 4-13:
Computer Management is your one-stop shop for system management utilities.



Computer Management is a Microsoft Management Console (MMC) snap-in that acts as a container for other MMC snap-ins. Its only offering or benefit over the individual tools is that you have to add only one snap-in to the MMC, rather than a whole series.

By putting all these utilities together in one location, system management is made substantially easier for most users.

In addition to using Computer Management to view and manage your computer, you can manage a remote computer that is someplace else on your network. To do this, right-click Computer Management (Local) at the top of the tree in the left-hand pane and choose Connect to Another Computer. This will open the Select Computer dialog box, which allows you to type the name of another computer on the network, and then click OK. Once connected to that other computer, you will be able to use all of the management tools, with the exception of the Disk Defragmenter, which will only operate on the local computer.



Lab 4-3 will walk you through the process of creating a custom Computer Management MMC. Lab 4-3 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Getting a Handle on Diagnostic Utilities

To figure out what's wrong with a faulty computer, you need to first find out what it is doing, both right and wrong. The tools in the following sections are designed to provide you with information on a variety of areas of your computer.

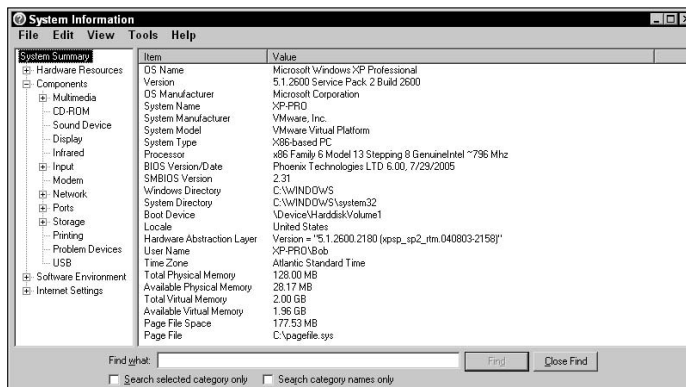
The System Information Tool

The System Information Tool (`msinfo32.exe`) is available in both the Windows 2000 and Windows XP operating systems. This tool was designed to provide information about as many components in the operating system as possible. Figure 4-14 should give you a good feel for the type of information that is available within this tool. If it's part of the operating system, then the System Information Tool can give you information about it.

To open the System Information Tool, choose Start→All Programs→Accessories→System Tools→System Information Tools.

Microsoft use to use a tool called `WinMSD.exe`, which was used to show similar information to what is displayed in the System Information Tool, which is now the latest iteration of a consolidation tool for system information. If you spend a few minutes investigating the data found in this tool, you will get a feel for the very detailed information it can provide. Very few areas of your system's hardware are not covered by this tool. It is well worth your time to get to know this tool to see what information it can gather.

Figure 4-14:
The System Information Tool provides information about your installation and hardware configuration.



If you want to save the current system information for archiving or to use as a baseline, you can choose File⇒Save to save the information into an `.nfo` file, which can later be opened and viewed with the System Information Tool.

In addition to the information that this tool provides, the Tools menu gives you links to other common tools that you might want to run, such as Net Diagnostics, System Restore, File Signature Verification Utility, DirectX Diagnostics Tool, and Dr. Watson.

Event Viewer

On some computers, you can check log files to look for errors with the system, applications, or hardware. Microsoft created a central logging and reporting tool that many components in the operating system and many applications use. Any application that wants to can log events to the event logs or even create its own event logs to be viewed with the Event Viewer. The Event Viewer is the tool that you use to view all of this logged information, as well as to configure log settings or clear the logs.

To view information, you open Event Viewer from your Administrative Tools using Start⇒Control Panels⇒Administrative Tools⇒Event Viewer. The interface has the log files on the left of the screen and the log data on the right. You will notice as you select the different logs on the left that there are five different log entry types and, by default, three logs. The three default logs are System, Application, and Security. The Security log has two event types, either a lock or a key for Failures or Successes. The other two logs have three different entry types — Information, Warning, and Error — and they are identifiable by the symbol or color (blue, yellow, and red, respectively). If you want to view log files on another computer, right-click Event Viewer (Local) in the left pane, and choose Connect to Another Computer from the menu. This will open the Select Computer dialog box, where you can type another computer name and choose OK.

Figure 4-15 shows a typical System log. The Security log has audit information about your system if you have enabled auditing of system events or file system access. The System log records major system events, like service failures, system startup or shutdown, and Stop events. The Application log records events from applications, although you may find that some applications record their events in the System log, and some items thought of as system events are recorded in the Application log.

Figure 4-15:
The System log has several types of messages telling you about major system events.

Type	Date	Time	Source	Category	Event	User	Computer
Warning	8/10/2006	10:32:28 AM	W32Time	None	14	N/A	XP-PRD
Error	8/10/2006	10:32:18 AM	Dhcp	None	1000	N/A	XP-PRD
Warning	8/6/2006	10:24:27 PM	MxSmb	None	3019	N/A	XP-PRD
Warning	8/6/2006	10:24:18 PM	MxSmb	None	3019	N/A	XP-PRD
Warning	8/6/2006	10:24:15 PM	MxSmb	None	3019	N/A	XP-PRD
Warning	8/6/2006	10:24:15 PM	MxSmb	None	3019	N/A	XP-PRD
Information	8/6/2006	10:18:21 PM	Service Control Manager	None	7036	N/A	XP-PRD
Information	8/6/2006	10:08:32 PM	Service Control Manager	None	7036	N/A	XP-PRD
Information	8/6/2006	10:08:31 PM	Service Control Manager	None	7035	SYSTEM	XP-PRD
Information	8/6/2006	10:08:16 PM	Service Control Manager	None	7036	N/A	XP-PRD
Information	8/6/2006	10:08:16 PM	Service Control Manager	None	7035	SYSTEM	XP-PRD
Warning	8/6/2006	9:37:54 PM	LsaSrv	SPNEGO ...	40961	N/A	XP-PRD
Information	8/6/2006	9:37:54 PM	Service Control Manager	None	7036	N/A	XP-PRD
Information	8/6/2006	9:37:54 PM	Service Control Manager	None	7035	SYSTEM	XP-PRD
Information	8/6/2006	9:37:54 PM	Service Control Manager	None	7036	N/A	XP-PRD
Information	8/6/2006	9:37:54 PM	Service Control Manager	None	7035	SYSTEM	XP-PRD
Information	8/6/2006	9:37:54 PM	Service Control Manager	None	7036	N/A	XP-PRD
Information	8/6/2006	9:37:54 PM	Service Control Manager	None	7035	SYSTEM	XP-PRD
Information	8/6/2006	9:37:54 PM	Service Control Manager	None	7036	N/A	XP-PRD
Information	8/6/2006	9:37:54 PM	Service Control Manager	None	7035	SYSTEM	XP-PRD
Information	8/6/2006	9:31:41 PM	W32Time	None	35	N/A	XP-PRD
Information	8/6/2006	9:31:53 PM	eventlog	None	6005	N/A	XP-PRD
Information	8/6/2006	9:31:53 PM	eventlog	None	6009	N/A	XP-PRD
Information	8/6/2006	9:31:30 PM	eventlog	None	6006	N/A	XP-PRD
Information	8/6/2006	9:29:37 PM	Distributed Link Trackin...	None	12505	N/A	XP-PRD
Information	8/6/2006	9:29:31 PM	Workstation	None	3260	N/A	XP-PRD
Information	8/6/2006	9:29:29 PM	Service Control Manager	None	7036	N/A	XP-PRD
Information	8/6/2006	9:29:29 PM	Distributed Link Trackin...	None	12505	N/A	XP-PRD

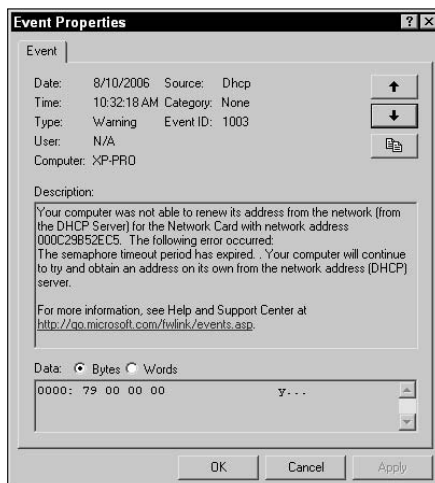
To see detailed information about the log entries, you can double-click any line item. When you do this, you see something that resembles Figure 4-16, listing date and time, source (system area or application), type of event, event ID (the error code returned by the application), user (if it is related to a user account), computer (if it is related to a computer account), a text-based description of the error, and data (which may show raw data related to the event as hexadecimal and ASCII data).



Read the Description section carefully because many developers put troubleshooting steps right into their error events.

If you want to navigate to the next or previous events, you can use the arrows in the top-right corner of the dialog box, and you can use the button with the sheet of paper on it to copy all of the event information to the Clipboard. You can then directly paste this information into a document or e-mail when reporting this event to other support people.

Figure 4-16:
The detailed information for an event gives you several relevant pieces of information.



If you can't decipher the event message, you can find additional help at www.eventid.net, which allows you to plug in the EventID and the source, and then you are able to view recorded events and descriptions of both the errors and the fixes. This site also contains links to Microsoft's Knowledge Base for articles related to the errors. In order to follow the links, you need to register on the EventID site, but because the site provides the Q article numbers (Microsoft's knowledge base has all articles number, with a Q as the initial character in the filename), you can go directly to <http://support.microsoft.com> and locate the articles yourself.

If you right-click a log file, you will have the option to save the event log as a text file or an `.evt` file, which can be loaded back into Event Viewer, or to clear or empty the event log. The View submenu has options to allow you to find specific entries or filter the log file to show certain events. You can filter the log by any of the field values.

If you want to change how long the log files are kept, you can right-click the log file that you want to modify and choose Properties. The default log file size is 512KB but can be set all the way up to 4194240KB, or 4GB, and there are settings on what to do when the log reaches its maximum size. When it reaches its maximum size, you can overwrite events as needed, overwrite events older than a number of days (default is 7 days), or not clear the events at all, in which case you stop logging new events.

In most cases, you should use Overwrite Events Older Than X Days and choose a number of days for which normal activity will not fill the log file. This system means that you can ignore the logging activity, but you will be

notified when the log file fills up, and you can start your investigation into the activity that filled it prematurely. I typically leave this setting between 7 and 14 days and make the size of the log file larger if necessary.

Network Diagnostics

The Network Diagnostics tool tests and reports on overall system health, with specific focus on the network. This tool is part the Windows XP Help and Support Center, but the easiest way to access it is through the System Information tool; therefore, you use Start⇨All Programs⇨Accessories⇨System Tools⇨System Information, and when the System Information window opens, choose Tools⇨Net Diagnostics. After you run a scan of your computer, the reported results are broken into three categories:

- ◆ **The Internet Services section** reports on Outlook Express and Internet Explorer settings.
- ◆ **The Computer Information section** reports on the same general information that is reported via the System Information Tool.
- ◆ **The Modems and Network Adapters section** reports on the settings and status of your modems, network adapters, and network clients after a series of tests.

If you want to save a copy of the findings, click the Save to File button, and the report will be saved to both your desktop and `C:\Windows\pchealth\helpctr\system\netdiag\Netdiag <Date> <Time>.htm`. This is just another tool that you can use to gather information to diagnose problems.

Dr. Watson

Dr. Watson is used to gather information about application crashes, which can be used by application developers to help diagnose errors, so you will use this information only when you're working with an application developer. To launch Dr. Watson, choose Start⇨Run and type `drwtsn32.exe`. If you have been using another debugger for your applications, you can make Dr. Watson the default debugger by running `drwtsn32 -i`.

Within Dr. Watson, you can specify the locations for log files and crash dump files. *Crash dump files* are similar to the binary log files that are created during a Stop error, but they only capture the memory data that is related to the crashed application. The capture configuration for Dr. Watson includes the number of instructions and errors to save that defines the range of data around the error that is captured. Just like Stop errors, you have the option to save full or mini logs, and the choice here depends on which debugging tools the developer is planning to use. Full logs are larger and contain all of

the information that the developer could want captured, but if the developer is only using a small amount of that data, then the mini dump may include enough data to resolve the problems with the application. Mini dump files are much smaller in size, which will make them easier to send the developer.

A Dr. Watson error log has information in the following categories:

- ◆ **Application error statement:** Includes the date, application name, and exception number.
- ◆ **System information:** Related to the computer that was running the application.
- ◆ **Task list:** Includes all running processes, like the Task Manager would list.
- ◆ **Module list:** Includes the memory addresses of .dll files and drivers.
- ◆ **Memory and state dump information:** In a variety of formats for each referenced thread.

System File Checker

The System File Checker (*sfc.exe*) is part of the Windows File Protection interface and must be run by an Administrator. With the proliferation of applications and with software developers replacing stock Microsoft .dll files with custom versions, there is a real need for Windows File Protection and the System File Checker. DLL (Dynamic Link Library) files contain compiled application code that is the same type of code that makes up an application. By saving this code as a DLL, the developer makes the code to many applications, so that it can be reused, and reduces application development time. One dialog box that is found in a system DLL is the common Save As dialog box. Windows keeps the cached copies of these files for your use when there is corruption in the original files. To illustrate this feature, you can delete both `c:\windows\system32\freecell.exe` and `c:\windows\system32\cards.dll`. By the time you have the second file deleted, you should see that the first file has already been replaced. The missing copy of the file has been replaced by the copy that was stored in the cache in `%systemroot%\system32\dllcache`. If the file was not in the cache, then Windows would have prompted you for the OS CD.

When you run the System File Checker, you have a choice of switches that you can use:

- ◆ `/SCANNOW`. Immediately scans all files for incorrect versions and verifies the versions that are stored in the cache.
- ◆ `/SCANONCE`. Scans all files on the next boot.
- ◆ `/SCANBOOT`. Scans all files at every boot.

- ◆ /REVERT. Changes `sfc.exe` to its default settings.
- ◆ /PURGECACHE. Purges the file cache to allow it to be rebuilt.
- ◆ /CACHESIZE=*x*. Sets the size, in megabytes, of the file cache.

Miscellaneous Utilities

A number of miscellaneous utilities are also available within Windows. The following sections take a look at some of them.

Task Scheduler

The system Task Scheduler service is responsible for running scheduled tasks. This service, like all services, is running in the background, and in this case, making sure that your tasks are performed at the times you have planned.

Your interface to this service is through the Start⇨All Programs⇨Accessories⇨System Tools⇨Scheduled Tasks, which opens the Scheduled Tasks folder.

To schedule a task, follow these steps:

1. Double-click Add Schedule Task.

The Scheduled Task Wizard will open.

2. Click Next to continue.

You will be presented with a list of applications that may be chosen and a Browse button to select applications that are not listed.

3. Select an application from the list or use the Browse button to locate an application you want to execute and then click Next.

4. Type a name for the task and select the scheduled interval for the task, which includes:

- Daily
- Weekly
- Monthly
- One time only
- When my computer starts
- When I log on

5. Once you have set a name and a scheduled interval, click Next.

Depending on your schedule interval, you will be presented with different scheduling times.

6. Specify the exact times you want the task to run, and when you are done, click Next.**7. Type the user credentials that you want the Task Scheduler to use when running the task — the user credentials will require appropriate rights to perform the actions you want performed — then click Next.**

You are now on the last screen of the wizard. You can choose to view the Advanced properties for the task, which include the settings we have already seen in the wizard, as well as these:

- Do not run the task if you have a laptop running on batteries.
- Maximum duration for the task to run.
- Run the task only if the computer is idle, and stop the job when the computer is no longer idle.

8. If you do not want to view these settings, just click Finish; otherwise, select the Open Advanced Properties for the Task When I click Finish Checkbox, and click Finish.

When you exit the wizard, you will see the new task in the Scheduled Tasks folder. To see the properties for any task, select the task and right-click → Properties. The properties of the task will allow you to adjust the command-line options for the application that is scheduled to run and change the schedule or the advanced options that were mentioned during the wizard. If you do not want to wait for the schedule, but rather have the task run immediately, then select the task, right-click it and select Run. You will see the status of the job change to Running. The other information that is shown in the Scheduled Tasks folder includes:

- ◆ When the task is scheduled to run next.
- ◆ When the task last ran.
- ◆ What the status was or what error code was returned the last time it was run.
- ◆ 0x0 is the standard code for a successful execution.

Windows Script Host

Windows Script Host (`wscript.exe` and `cscript.exe`) is a utility that doesn't fit easily into any other categories. Windows Script Host is used to execute Visual Basic Script (VBScript) or JavaScript. The tasks that you are

able to accomplish by using VBScript are numerous and almost unlimited. Most Windows-based applications support ActiveX automation and can be controlled through VBScript. You can use Windows Script Host to execute such scripts.

Windows Script Host has two components that can execute script: `wscript.exe` and `cscript.exe`. The `wscript.exe` program executes VBScript from within the Windows environment; `cscript.exe` executes scripts from the command prompt. Both these utilities are interpreters for VBScript. If the scripts were converted into true Visual Basic and compiled into executables, you would achieve much better performance.

Windows Script Host was introduced to the OS with Windows 98 and is included with all current versions of Windows. To get code samples or information about using Windows Script Host, visit <http://msdn.microsoft.com/scripting>. You can also get scripting information from the Hey, Scripting Guy! site through www.microsoft.com/technet/scriptcenter/learnit.mspx or www.microsoft.com/technet/scriptcenter/default.mspx.



Lab 4-4 has you create a few simple scripts using Windows Script Host. Lab 4-4 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

edit.com

The `edit.com` program allows you to modify the contents of text files. We cover `edit.com` in detail in Book V, Chapter 5. The version of `edit.com` that ships with Windows XP includes some enhancements over the original version that ships with MS-DOS. The two largest enhancements are:

- ◆ **The ability to open multiple files**
- ◆ **The ability to use standard Windows shortcut keys for Cut, Copy, and Paste operations**

expand.exe

You can use the command line `expand.exe` utility to expand individual source files that are shipped in a compressed format on your OS CD. This utility can be used to replace corrupt files that are being used by the operating system.

Files that have been compressed for expansion with `expand.exe` are found on your OS CD with the last letter of their filenames replaced with an underscore, or they are compressed as `.cab` files.

Common tasks you might use `expand` include:

- ◆ Viewing the names of all the files in a `.cab` file, such as seeing the names of all of the `.gif` files contained in `iis6.cab` file on the Windows XP Installation CD, where `D:` is the drive letter of the CD drive that contains the Windows XP installation CD:

```
expand.exe -D D:\i386\iis6.cab -F:*.gif
```

- ◆ Extracting a file from a `.cab` file, such as extracting the `IIS_winxp.gif` file from the `iis6.cab` file to the root of the `C:` drive:

```
expand.exe D:\i386\iis6.cab -F:IIS_winxp.gif C:\
```

- ◆ Expanding a compressed file — for example, extracting a fresh copy of the Freecell Help file to the root of the `C:` drive:

```
expand.exe D:\i386\freecell.ch_ c:\freecell.chm
```



Lab 4-5 will give you some practice using this command line. Lab 4-5 can be found in the `Labs.pdf` file in the `Author` directory of the CD-ROM.

Getting an A+

This chapter examines several Windows-based utilities. You find out about the following:

- ◆ Disk maintenance programs, such as `fdisk`, `defrag`, and `chkdsk`, are used to improve performance and reliability of the disk subsystem.
- ◆ Configuration utilities, such as the System Configuration Utility, System Information Tool, Registry editors, Registry scanners, Device Manager, and the Computer Management console, are used view or change Windows configuration settings.
- ◆ `edit.com` can be used to create or modify text files when `notepad.exe` or other editors are not available or appropriate.
- ◆ `expand.exe` can be used to decompress Windows XP source files.

Prep Test

- 1 defrag.exe does what for your computer?**
 - A Rearranges memory so that access to it is improved.
 - B Removes dust and fragments from your computer.
 - C Rearranges data on your disk drive so that access to it is improved.
 - D Recovers files that have been corrupted.

- 2 chkdsk.exe can do which of the following?**
 - A Rearrange data on your drive so that disk access is faster.
 - B Consolidate free space on your hard drive.
 - C Correct problems with file storage on your hard drive.
 - D Change the partition table of the drive from FAT16 to FAT32.

- 3 You have to edit several of the configuration files on your computer. Which of the following utilities might you use?**
 - A systmedt.exe
 - B editcnfg.exe
 - C msconfig.exe
 - D setfile.bat

- 4 What programs can be used to edit the Registry on a Windows XP computer?**
 - A registry.exe
 - B cfgedit.exe
 - C regedt32.exe
 - D regedit.exe

- 5 What command is used to export a section of your Registry while using regedit.exe?**
 - A Registry⇨Export Registry File
 - B Registry⇨Extract to Recovery File
 - C Export⇨Set Restore Options
 - D Registry⇨Backup Key

- 6 What tool should you use to change an exported Registry file?**
 - A sysedit.exe
 - B regedit.exe
 - C edit.com
 - D Microsoft Word

- 7** Which files on your hard drive make up the Windows XP Registry?
- A system.dat and registry.dat
 - B user.dat and system.dat
 - C user.dat and hardware.dat
 - D system, sam, security, and ntuser.dat
- 8** Which operating system uses `regedt32.exe`?
- A MS-DOS
 - B Windows 95
 - C Windows 98
 - D Windows 2000
- 9** You need to find out which memory addresses are being used by a driver. What tool will tell you what you need to know?
- A `memedit.exe`
 - B System Resource Meter
 - C Device Manager
 - D System Driver Checker
- 10** You need to automate a procedure on your computer. What tools could you use?
- A `wscript.exe`
 - B `ScriptIt.exe`
 - C `config.sys`
 - D Local ASP
- 11** You need to automate the process of copying information from a Microsoft Excel spreadsheet into a Microsoft Word document. What tools could you use?
- A `wscript.exe`
 - B `ScriptIt.exe`
 - C `config.sys`
 - D Local ASP
- 12** What is the name of the command-line version of Windows Script Host?
- A `ComScript.exe`
 - B `wscript.exe`
 - C `cmscript.exe`
 - D `cscript.exe`

- 13** You need to create a new Visual Basic script for Windows Script Host. What tool will you use?
- A Script Developer
 - B Script Editor
 - C Notepad
 - D Visual Basic Script IDE
- 14** You need to replace `notepad.exe`, which was accidentally deleted from your Windows XP computer. What command can you use to retrieve a copy from the original `.cab` files?
- A `extract.exe`
 - B `expand.exe`
 - C `filerepl.exe`
 - D `cabget.exe`
- 15** What command is used to create or modify text files?
- A `textedit.com`
 - B `notepad.com`
 - C `edit.com`
 - D `fileedit.com`
- 16** What command is used to change the partitioning on your hard drive?
- A `format.com`
 - B `fdisk.exe`
 - C `diskpart.exe`
 - D `repart.exe`

Answers

- 1** C. `defrag.exe` rearranges data on your disk so that access to files is faster. By allowing your disk to perform reads of data in contiguous blocks, it takes less time to read files. See *"fdisk.exe and Disk Management."*
- 2** C. `chkdsk.exe` fixes minor problems with files that are stored on your hard drive and can also test and verify the clusters on your drive. Review *"chkdsk.exe (Check Disk)."*
- 3** C. `msconfig.exe` allows you to edit several configuration files at the same time. These files include `boot.ini`, `win.ini`, and `system.ini`. Check out *"System Configuration Utility."*
- 4** D. `regedit.exe` is the editor for Windows XP computers. If you are using Windows 2000, then you can use either `regedit.exe` or `regedt32.exe`, but with Windows XP, if you run `regedt32.exe`, you will actually open `regedit.exe`. Peruse *"regedit.exe."*
- 5** A. Export Registry File from the Registry menu will export a section of your Registry. Take a look at *"regedit.exe."*
- 6** C. Registry import files must be straight text. Notepad and many other programs that edit text files add different characters to handle carriage returns and line feeds. This can create files that will not be imported through `regedit.exe`. Peek at *"edit.com."*
- 7** D. `system`, `sam`, `security`, and `ntuser.dat`, in addition to software, are the files that make up the Registry. All of the files are found in `<%systemroot%>\system32\config`, with the exception of `ntuser.dat`, which is found in each user's profile directory. Look over *"regedit.exe."*
- 8** D. Only Windows NT, Windows 2000, and Windows XP use `regedt32.exe`. Study *"regedt32.exe."*
- 9** C. Device Manager will tell you what resources are being used by which devices. Refer to *"Device Manager."*
- 10** A. `wscript.exe`, or Windows Script Host, can be used to automate processes. Examine *"Windows Script Host."*
- 11** A. `wscript.exe` can control applications that support ActiveX automation, such as Microsoft Word and Microsoft Excel. See *"Windows Script Host."*
- 12** D. `cscript.exe` is the command-line script interpreter. Review *"Windows Script Host."*
- 13** C. Notepad or another text editor can be used to create or edit your Visual Basic scripts. Check out *"Windows Host Script."*

Using Windows-Based Troubleshooting Utilities

- 14 B. `expand.exe` is used to extract files from the source `.cab` files, while `expand.exe` can also be used to uncompress the individual source files for Windows XP. *Peruse "expand.exe."*
- 15 C. `edit.com` can be used to create or modify text files. *Take a look at "edit.com."*
- 16 B. `fdisk.exe` is used to change the partition table on your hard drive. *Peek at "fdisk.exe and Disk Management."*

Book VII

Recovering Systems

The 5th Wave

By Rich Tennant



“Drive carefully, remember your lunch, and always make a backup of your directory tree before modifying your hard disk partition file.”

Contents at a Glance

Chapter 1: Managing Error Codes and Startup Messages	793
Chapter 2: Managing Common Problems	811
Chapter 3: Preparing for Disasters with Disaster Recovery.....	831

Chapter 1: Managing Error Codes and Startup Messages

Exam Objectives

- ✓ Recognizing and resolving common error codes and startup messages
- ✓ Identifying steps to correct operational problems
- ✓ Using diagnostic utilities and tools to resolve operational problems

This chapter examines problems that are common to the Windows environment. Many of these problems occur during the boot process, but a few can occur at any time. I examine boot or startup errors first and then other errors that occur within the operating system. The end of the chapter reviews some tools that are used to help diagnose the cause of the errors.

As an A+ Certified Professional, you will often need to diagnose problems based on error codes and messages from the operating system. This chapter introduces you to many of the common boot or startup error messages that you may see when using Windows. You will also be introduced to the key files that are required to allow you to boot a computer running Windows.

MS-DOS and Windows 9x Boot Errors

The Windows 9x and MS-DOS boot processes are identical, up until the point that the Windows 9x GUI is started with the execution of `win.com`. This section will look at the errors that you may encounter during the MS-DOS portion of the boot process, prior to `win.com` starting. In many cases, because of a need to support FAT32 or large disks, many people use a Windows 98 boot disk instead of an MS-DOS 6.22 boot disk.

A Windows 9x boot disk is used regularly as a rescue medium or during a pre-OS installation phase so it is extremely useful to know and be able to troubleshoot the boot process of this disk. Since I am not actually booting Windows with disk, but rather the MS-DOS command shell that ships with Windows 9x, I will be referring to this disk as an MS-DOS boot disk. Most of the problems surrounding the boot process involve the five main startup

files: `io.sys`, `msdos.sys`, `config.sys`, `autoexec.bat`, and `command.com`. The following sections look at some of the common error messages you may receive.

Error in config.sys line XX

When working with an MS-DOS boot disk, your `config.sys` and `autoexec.bat` files, if they exist, are processed during startup. Depending on the settings found in your `msdos.sys` file, there may or may not be an animated Windows 9x logo. The purpose of the Windows 9x animated logo is to hide the processing of these two files. When processing the `config.sys` file, any errors in the file are reported to the screen, and the animated logo is suspended.

`config.sys` and `autoexec.bat` are required startup files for MS-DOS but are used by Windows 9x to override default settings that Windows 9x already uses or to allow backward compatibility with MS-DOS-based applications. If you do not need to override settings or provide backward compatibility, your system doesn't need these two files. For more information about these files and the Windows 9x boot process, consult Book V, Chapter 6.

If you want to see all of the lines in the `config.sys` process, press the Esc key when you see the text "Starting Windows 9x" or anytime after the animated logo is on the screen. If you press the Esc key after the animated logo is displayed, the animated logo will show up again when the monitor's refresh rate is adjusted, so you will have to press the Esc key a second time to remove it again.

If there is a problem with the `config.sys` file, you may see an error message that resembles this section heading, *Error in config.sys line XX*. As the error message suggests, there is an error in line `xx` of your `config.sys` file. `config.sys` is a text file on the root of your boot drive. Because it's a text file, you can open it with any text editor, such as `edit.com` or `notepad.exe`, and change it. Count down the number of lines indicated by the error message, starting with one, to get to the offending line. This line likely has some type of syntax error. Check to make sure the path or file that is referenced by the line in `config.sys` matches the actual location. This error can also result from passing the wrong parameters to a program or driver.

Bad or missing command.com

`command.com` represents the 16-bit real-mode component of the Windows 9x OS. It is called extremely early in the boot process — immediately after the boot loader `io.sys` configures the environment by using `msdos.sys` and `config.sys`. Because the real-mode environment actually loads the 32-bit

graphical environment, a missing or corrupt `command.com` results in a computer that will not boot. In order to replace or repair `command.com`, you need a working copy of `command.com`. This can be found on any other Windows 9x boot disk with the same version. If you don't have another boot disk, you can create one by using the Startup Disk tab of the Add/Remove Programs control panel of any Windows 9x computer. You may also make a bootable disk on another Windows 9x computer by formatting a disk and including the system files. If you do not have a working installation of Windows 9x, then you may find images of these disks on the Internet. Ed has had good luck downloading disk images from www.bootdisk.com.



Always use care when downloading files from Web sites that you are not familiar with.



In Book VII, Chapter 3, we discuss how each different version of Windows 9x has a different version of the Startup disk. When replacing system files, you must use the correct version of the disk for your version of Windows.

If you have a valid, current Startup disk, you can use it to boot your computer to a command prompt. Then type this:

```
sys c:
```

This copies a fresh set of boot files over to drive `c:`. These files include `io.sys`, `msdos.sys`, and `command.com`.

No operating system found

Although the “No operating system found” error can affect MS-DOS, Windows 9x, and Windows NT–based computers like Windows XP, this section deals with solutions for MS-DOS and Windows 9x. For solutions from a Windows XP perspective, read the next section, “Windows XP Boot Files and Boot Errors.”

When dealing with MS-DOS and Windows 9x, this error message means that there is actually a problem with the boot loader program in the boot sector of your drive. Running `sys.com` on your boot drive should solve the problem. There is also a chance that a formatted floppy disk was left in the floppy drive and it does not have a boot loader program.

Windows XP Boot Files and Boot Errors

Just as MS-DOS and Windows 9x suffer from boot problems, so does Windows XP. This section takes a look at some of the boot problems that

affect Windows XP computers. These problems will affect all Windows NT-based computers, such as Windows 2000, Windows Server 2003, and Windows Vista.

The Windows XP boot process may suffer from boot sector corruption, boot loader problems, and drive identification problems within `boot.ini`. For most problems, the Emergency Repair Process or Recovery Console are reasonable troubleshooting and repair steps. (I discuss the Emergency Repair Process and the Recovery Console in Chapter 3 of this minibook.)

SCSI issues

The `boot.ini` file on your hard drive identifies the controller bus that contains the hard drives in your computer. This controller bus is identified as either SCSI or multi. If your drives are connected to a SCSI controller with the BIOS disabled, the controller is identified as SCSI. In all other circumstances, the controller is identified as multi. If your system is unable to initiate the boot process, or if you receive an error regarding a missing `ntoskrnl.exe`, then the problem could be a misidentified SCSI controller. Book V, Chapter 6, has more information about the `boot.ini` file.

No operating system found

The “No operating system found” error in Windows XP, just like in MS-DOS, is tied to severe corruption of the boot sector on your hard drive. To restore your drive to working condition, you should perform the Emergency Repair Process to restore the boot sector on your hard drive. The Emergency Repair Process is covered in Book VII, Chapter 3. Minor corruption of the boot files would report missing files, such as `ntldr` or `ntdetect.com`, and the next section will tell you how to deal with missing files. As with MS-DOS, the “No operating system found” error could be caused by having left a formatted floppy disk with no boot loader in the floppy drive during boot.

Missing boot files

Windows XP requires four core files to boot, and a fifth file is optional. The required files are

- ◆ `ntldr`
- ◆ `ntdetect.com`
- ◆ `boot.ini`
- ◆ `ntoskrnl.exe`

The optional file is

- ◆ `ntbootdd.sys`, which is the SCSI adapter driver that Windows XP boots from if the SCSI adapter doesn't have an active BIOS.

In the following sections, you get a look at what you can do if any of these files are missing or appear to be missing on your computer. For further information about the Windows XP boot processes, consult Book V, Chapter 6.

ntldr

If `ntldr` is missing, you will receive the following error message when your computer tries to boot:

```
NTLDR is missing
Press any key to restart
```

This message often means that a formatted floppy was left in the disk drive. If this file is actually missing, it can be replaced with a working copy from any Windows NT computer; although you should try to replace the missing file with a copy from the same or a newer version of Windows to maintain full compatibility. To replace this file, you may have to make a boot disk to get your computer booted. This can be done with the following steps:

- 1. Format a disk on a working Windows XP computer.**
- 2. Copy the files `ntldr`, `ntdetect.com`, and `boot.ini` from the root of the C: drive.**

These files are hidden, system, and read-only, so you will have to modify their attributes to copy them. Check out Book V, Chapter 4, for the low-down on modifying attributes.

- 3. Edit the `boot.ini` file to reflect the boot configuration of your target computer (the one that will not boot).**

boot.ini

If the `boot.ini` file is missing and you have installed Windows XP in its default location, then the system will boot but will not display a boot menu. The default location for Windows XP and newer versions of Windows is `c:\windows`; for Windows 2000 and Windows NT, the default location is `c:\winnt`. If you have not installed Windows in its default location, you will receive a message stating that `ntoskrnl.exe` or `hal.dll` is corrupted or missing and that you should replace the file. The message will look similar to this one from Windows 2000:

```
Windows 2000 could not start because the following file is missing or corrupt:
<Windows 2000 root drive>\system32\ntoskrnl.exe
```

The message may also be similar to this one from Windows XP:

```
Windows could not start because the following file is missing or corrupt:  
<Windows root>\system32\hal.dll  
Please re-install a copy of the above file.
```

This message is misleading because the problem is really with the `boot.ini` file. The reason for the message is that the boot loader (`ntldr`) has gone to the default location, and `ntoskrnl.exe` or `hal.dll` files were not there. If the `boot.ini` file is replaced and the boot path is correct for your installation, then the boot process will continue as normal.

`boot.ini` is a text file and can be edited with any text editor, such as `notepad.exe`. To replace the `boot.ini` file, you may need to create a boot disk as you did to replace the missing `ntldr` file.

ntdetect.com

If `ntdetect.com` is missing, you will receive the following error message:

```
NTDETECT failed
```

This file is generic, like the `ntldr` file, and can be replaced in the same way that you replace the `ntldr` file. See the section, “`ntldr`,” earlier in this chapter.

ntoskrnl.exe or hal.dll

You should be able to find the `ntoskrnl.exe` file in the `<Windows XP root drive>\windows\system32\` directory. If it’s missing, you receive a message like this message from Windows 2000:

```
Windows 2000 could not start because the following file is missing or corrupt:  
<Windows 2000 root>\system32\ntoskrnl.exe
```

The message may also look like this one from Windows XP:

```
Windows could not start because the following file is missing or corrupt:  
<Windows root>\system32\hal.dll  
Please re-install a copy of the above file.
```

These messages may look familiar; they are the same messages that you receive if your `boot.ini` file is misconfigured. After checking the `boot.ini` file, if the `ntoskrnl.exe` or `hal.dll` files are actually missing, you will have to do one of the following:

- ◆ Perform an Emergency Repair Process for Windows NT 4.0 or Windows 2000 to replace any missing or corrupted files on your system.

For more information on Emergency Repair Process, see Book VII, Chapter 3.

- ◆ Attempt to replace the file by using the Recovery Console.
- ◆ Re-install the operating system.



Ninety-nine out of 100 times, the missing `ntoskrnl.exe` or `hal.dll` error message means an error with the `boot.ini` file, and the `ntoskrnl.exe` or `hal.dll` file is fine.

ntbootdd.sys

Most systems boot either from IDE/ATA hard drives or from SCSI drives that are attached to a SCSI controller or adapter that has a working BIOS. In either of these cases, you won't find the `ntbootdd.sys` file on your drive. If you are booting from a SCSI drive that is attached to a SCSI controller that has its BIOS disabled, then you will find the `ntbootdd.sys` file on the root of your bootable drive, along with `ntldr`, `ntdetect.com`, and `boot.ini`. This file is computer-specific because it's the SCSI driver for *your* SCSI controller. It can be replaced with a copy of the file found on any other computer that boots from the same SCSI controller. If you don't have another system with the same configuration, you can get a copy of the SCSI driver from either the driver disk for the SCSI controller or possibly from your Windows installation CD-ROM. This driver will have to be renamed to `ntbootdd.sys` and copied to the boot drive, and you may need to make a boot disk like the one described in the section "ntldr," earlier in this chapter.

Device Related Errors

It is unfortunate that the devices and their drivers that allow us to accomplish so much of our day-to-day work with computers are also one of the biggest factors in not being able to do work on our computers. Ideally, when all the devices are configured on your computer, you should be able to work with no problems from your drivers. Most people's computers don't remain in a static mode but rather are in constant flux. Even though devices are working fine, many people feel the need to try to improve performance by changing settings, upgrading drivers, or installing service packs. Although upgrading drivers and installing items such as service packs are common practice, they should be done carefully. A service pack, for instance, can change the way all of the drivers on your computer function. In the rare case in which something does go wrong, you may find that the fix is difficult, but in most cases, it will be related to a file version or configuration setting. This section will take a look at how to address these problems.

A device referenced in system.ini, win.ini, or Registry is not found

From time to time, you will find that one of your startup files still references a device that you thought had been removed from your system. The files that reference devices include `system.ini`, `win.ini`, and the Registry. If this happens, you may have to edit these files manually in order to fix the problem. If an error message tells you that a referenced device does not exist, then take note of the device that is being referenced because you will have to search for it in your startup files. Most devices are found within your `system.ini` file, so this is the first file that should be checked. Both `ini` files can be opened and edited with `notepad.exe` or `sysedit.exe`. To find out how to use `sysedit.exe`, take a look at Book V, Chapter 6.

If the device is listed in the Registry, then it should be listed in Device Manager. You open the Device manager by choosing Start→Control Panel→System, clicking the Hardware tab on the resulting dialog box, and then clicking the Device Manager button. Locate the device in Device Manager and delete it. If the device is still physically present in the computer, it will be re-added to Device Manager when your computer is rebooted. If you keep removing the device and it keeps coming back, that is because it is still physically present. Physically remove the device first, and then remove it from Device Manager.

In Windows XP, if you do not see the device that you want to remove in Device Manager, then you can select Show Hidden Devices from the View menu. If the device is not listed in Windows 2000, you can use the Add/Remove Hardware Wizard. If you can't find the device in the Windows GUI, then you can attempt to search the Registry to locate the device and correct the issue. To find out about editing the Registry, see Book VI, Chapter 4.

Registry corruption

There are two main ways that the system Registry can become corrupted: Updates to the Registry via one of the Registry editing tools or an import of a registry settings file, and by the files that make up the Registry becoming damaged or deleted.

There are many ways to import data into the Registry, but most of them involve storing settings in a file and importing that file into the Registry. If the settings in the file are incorrect, you might be able to just continue computing without any problems, or you could end up with a system that no longer boots normally. If your system will not boot normally, then your only option would be to boot the system by using an alternative method, such as the Recovery Console, and replace the base Registry files with an untainted version.

You'll also need to use an alternative method to boot the system if the Registry files on the drive itself have become corrupted, and then you'll have to replace the Registry files. The user portion of the Registry is found in your user profile directory and is named `ntuser.dat`. Your user profile directory is in `C:\Documents and Settings\`. The system portion of the registry is found in `%systemroot%\system32\config`, in the files `SAM`, `SECURITY`, `system`, and `software`. For information about editing the Registry, see Book VI, Chapter 4.

Safe Mode

In order to let you repair the operating system from within the operating system, Microsoft has provided Safe Mode. *Safe Mode* is available with Windows XP, Windows Server 2003, Windows Vista, and Windows 2000 and is a special boot of Windows that loads a minimal set of drivers. The only drivers that are loaded are the ones that are required to get the operating system running. Instead of loading the normal video driver, Safe Mode loads a 16-color VGA graphics driver. The `config.sys` and `autoexec.bat` files are skipped completely. If you have issues with drivers or driver configuration, then booting into Safe Mode can allow you to bypass these driver related problems, so that they can be fixed.

You can enter Safe Mode by pressing the F8 key when the operating system is booting up. If your computer boots into Safe Mode, the words "Safe Mode" appear in each corner of your desktop. If Windows XP fails to boot properly, then it will suggest, and attempt, to boot into Safe Mode on the next boot. If your computer boots into Safe Mode automatically, the last boot process was likely interrupted (usually by the user). For more information about Safe Mode and other boot methods, consult Book VII, Chapter 3.

Other Errors

As with all things in life, some things cannot easily be categorized, so this section discusses errors that do not fit in the other categories in this chapter.

Paging file or Swap file errors

The Windows *swap file* or *paging file* is a hard drive file that is used as additional RAM memory. Typically swap file is the name used to describe the file when working with Windows 9x, while with Windows XP the file is called the paging file or page file. The location of the paging file is recorded in the Registry. If the drive that contains the paging file becomes too full, you may encounter errors informing you of this fact. If this happens, create some additional space for the paging file by deleting some unnecessary files,

reducing the size of the paging file, or moving the paging file to a new location. The default location for the paging file is in the same drive as your Windows directory. Windows XP allows you to move the paging file to an alternate drive. If you have done so, and that drive has been removed from your system, you encounter errors telling you that the paging file could not be created. If this happens, configure Windows to use another drive for the paging file as shown in Book VI, Chapter 3.

Failure to start GUI

Sometimes, Windows won't be able to complete the boot process, but it won't generate an error, either. Instead, it will seem to hang at one spot in the boot process without going any further. You can deal with this problem in a number of ways.

The first action you might want to take is to see if the OS has loaded the networking components, which you can do using the `ping` utility which is discussed in Book VIII, Chapter 3. If the networking services are running, then you may be able to connect to the computer using Computer Management or Event Viewer, as described in Book VI, Chapter 4, which will allow you view information about the OS and read the event logs. There may be errors in the event log which will let you know what the problem is.

If you are not able to connect to the computer, then your only options are to continue to wait, or cycle the power on the computer. If the computer has locked up hard enough, then the soft power button may not work, and you may have to unplug the computer; if it is a laptop, you will have to remove the battery. This action, as with any power interruption, has the risk of causing corruption to the file system, and should be used as a last resort. On reboot, the system may boot properly, or you may want to use one of the F8 boot options which are discussed in Book VII, Chapter 3, such as Safe Mode. Safe Mode would allow you view the Event Logs using Event Viewer, and that may give you insight into the problem. The F8 boot options will also allow you to create a bootlog, which if your next boot is unsuccessful, will let you know approximately where the failure is occurred. If you are unable to find the solution, disconnect all peripherals such as USB devices, and attempt to reboot the computer. If the problem is with a peripheral or its driver, then you may see a successful boot. In some cases, I have seen errors with peripherals occur after the system has successfully booted for years, and the problem can usually be traced to an OS update, a drive update, or file system corruption of the driver.

If you think you know what is causing the problem, perhaps a driver or service that won't start, then you could also use the Recovery Console to try to fix those problems. For more information about the Recovery Console and its use, consult Book VI, Chapter 2.

Error Diagnostic Tools

The following sections examine the different tools you can use when diagnosing errors.

Dr. Watson

Dr. Watson, which was first included with Windows 3.x, has been designed to help troubleshoot problems on your system. It helps troubleshoot problems after they have occurred by generating log files and system snapshots.

For the Windows 9x operating system, Dr. Watson records extremely detailed information about what processes are currently running. This information is stored on the tabs that are located across the top of the Dr. Watson window. It breaks these processes into categories such as MS-DOS, 16-bit Windows, and 32-bit Windows. The Diagnosis tab allows the user to record additional information about what he or she was doing when the error occurred. Entering this information may be useful during the troubleshooting process. To load Dr. Watson for Windows 9x, use Start→ Run and type `drwatson.exe`. Dr. Watson then shows up in your system tray.

Figure 1-1 shows Dr. Watson's configuration screen for a Windows XP computer. Configure Dr. Watson in a Windows XP computer by running `drwtsn32.exe` from the Run command. Dr. Watson is capable of taking the entire contents of an applications memory space, and writing that information to a file on your disk, also called a crash dump file. This may be useful for the application developer to troubleshoot the error. Dr. Watson also logs and creates crash dump files for a variety of OS background services, which are really just applications, as shown in Figure 1-2. With Windows XP, Dr. Watson is always ready to generate these logs and crash dump files.

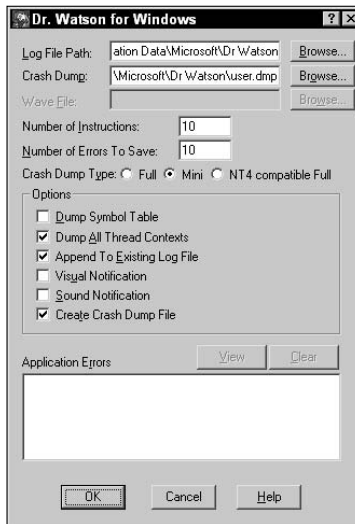
Windows Error Reporting

With Windows XP, Microsoft introduced a new tool for error reporting that has replaced most of the functionality originally covered by Dr. Watson. The new Error Reporting tool allows you various levels of reporting. By default, it is configured to report errors in both applications and the operating system directly to Microsoft, which allows its programmers the potential to use the information when creating new patches for the operating system. This “call home” functionality can be disabled by going to Error Reporting button on the Advanced tab of the System Control Panel applet. The Error Reporting button will open the Error Reporting dialog that is shown in Figure 1-3. This dialog allows you to disable the call home functionality but still have Windows report the errors to you; specify that Error Reporting works for both OS components and program; or use the Choose Programs button to specify which programs you want to have Error Reporting work with, or exclude.

Figure 1-1: Dr. Watson provides the Diagnosis tab for you to describe what you were doing when the error occurred.



Figure 1-2: How Dr. Watson deals with applications that crash.



When an application crashes you will see a message similar to the one in Figure 1-3; while if Windows has a Stop error, which is described in Book VII, Chapter 2, you will see the message after you reboot and log in. By clicking the link at the bottom of the dialog box, you may find some additional information related to the error. When you see this dialog, you have the option of

sending a summary of the error data and system state to Microsoft, as shown in Figure 1-4. If you disabled Error Reporting but kept the notify option enabled, you get a smaller dialog box which will not have an option for notifying Microsoft, and only an OK button to acknowledge the error.

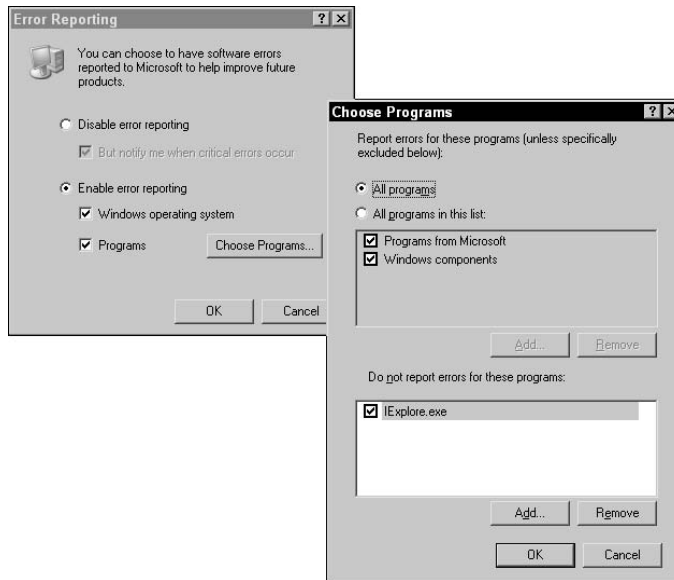


Figure 1-3: Windows Error Reporting helps Microsoft fix problems with the OS.




Figure 1-4: Even with Error Reporting enabled, you have the ability to control what errors go to Microsoft.

Event Viewer

Windows XP offers various logging tools, the greatest of which uses the Event Log service. The Event Log service logs errors and events into several different log files. The Event Viewer is the application you use to view the

contents of these log files. Windows XP always has at least three default logs: Application, Security, and System. These log files have a default size of 512KB each and automatically overwrite events after seven days. These settings can be adjusted for each file by right-clicking the log file in Event Viewer and choosing Properties to open the System Properties dialog box, shown in Figure 1-5.

Figure 1-5: The default settings for the log files should be enough to prevent filling.



In the event that any log fills up, you receive a pop-up message. When this happens, open the Event Viewer by choosing Start⇒Control Panels⇒Administrative Tools and then Event Viewer. To clear the log, right-click the log and choose Clear All Events. When you clear the log, you have the opportunity to save the events to a file.



Before clearing the filled log, examine the most recent events for an error that might have caused the log to fill quickly.

Getting an A+

This chapter examines a number of common errors that you are likely to encounter with MS-DOS, Windows 9x, or Windows XP computers. These errors range from configuration settings and files to boot files and driver-related problems. You also see some common resolution methods, such as booting into Safe Mode. Some key elements to remember in this chapter are:

- ◆ “No operating system found” is an MS-DOS error message, and may mean a floppy disk was left in your Windows XP computer.
- ◆ Most Windows XP boot sector or boot file errors can be fixed with either the Emergency Repair Process or the Recovery Console.
- ◆ Safe Mode is used to diagnose and resolve driver or startup problems.
- ◆ Event Viewer should always be checked as part of the troubleshooting process, to see what error messages are being reported by Windows.

Prep Test

- 1** Which of the following options are valid procedures for trying to resolve corrupted boot files with Windows XP? Choose all that apply.
- A Recovery Console
 - B `sys.com`
 - C `fdisk.exe`
 - D Emergency Repair Process
- 2** Which of the following are loaded when you boot into Safe Mode? Choose all that apply.
- A Your normal video driver.
 - B Your mouse driver.
 - C `config.sys`
 - D `autoexec.bat`
- 3** What is the purpose of a swap file?
- A To prepare files to be copied to other disks.
 - B To act as additional memory for the system.
 - C To act as an extension to the hard drive.
 - D To prepare files that are saved to disk.
- 4** An application has crashed, and you would like to know what else was running on the computer at the time because you think that this may be helpful in diagnosing the problem. Which program should you run?
- A `whatsup.exe`
 - B `sysrun.exe`
 - C Device Manager
 - D `drwatson.exe`
- 5** Where do you go to read and clear the Windows XP Event Log?
- A Event Log
 - B Event Viewer
 - C Log Reader
 - D Disk Cleanup

- 6** You have installed Windows 2000 in the default location, `C:\winnt`. What will be the effect of deleting your `boot.ini` file?
- A Your system will boot as normal, but will not display a boot menu.
 - B Your system will not boot and will display a “missing `boot.ini`” error message.
 - C Your system will not boot and will display a “missing `ntoskrnl32.exe`” error message.
 - D Your system will not boot and will display a missing “`ntdetect.com`” error message.
- 7** You are missing your copy of `ntldr`, and your Windows XP computer will not boot. How can you replace your copy of `ntldr`?
- A By using the `retrieve.exe` command to get a copy from the Windows XP installation CD-ROM.
 - B The only way you can replace this file is through the Recovery Console.
 - C This file is not required and does not have to be replaced.
 - D By copying it from any other computer that is running Windows XP.
- 8** You have just noticed that your `ntbootdd.sys` file is missing from your computer. What is the `ntbootdd.sys` file?
- A Text configuration information for your Windows XP system.
 - B The boot loader.
 - C The SCSI driver for your SCSI adapter if the BIOS is disabled on the adapter.
 - D The boot menu configuration file.
- 9** When should you boot into Safe Mode?
- A At every boot-up.
 - B When your system will not boot normally.
 - C When your computer is connected to the Internet.
 - D Only to play `freecell.exe`.
- 10** The system Registry for Windows XP is extremely important. What file or files make up the user portion of the Registry?
- A `registry.dat`
 - B `system.ini` and `registry.dat`
 - C `user.dat`
 - D `ntuser.dat`

Answers

- 1 A, D. Both the Recovery Console and the Emergency Repair Process allow you to repair problems with the boot files. *Check out "Windows XP Boot Files and Boot Errors."*
- 2 B. Your mouse driver is the only item on the list that loads when you enter Safe Mode. It also loads a VGA-compatible video driver and a minimal set of device drivers. *Peruse "Safe Mode."*
- 3 B. The Windows swap file acts as additional memory, or virtual memory, for the OS. *Take a look at "Paging file or Swap file errors."*
- 4 D. Dr Watson provides detailed information regarding what was running on the computer. *Peek at "Dr. Watson."*
- 5 B. Event Viewer is used to read, back up, and clear your event logs. *Look over "Event Viewer."*
- 6 A. Your system will boot as normal; when `boot.ini` is missing, Windows 2000 attempts to boot from the default installation location. *Study "Windows XP Boot Files and Boot Errors."*
- 7 D. `ntldr` is identical across all copies of Windows XP, so it can be copied from any working installation of Windows XP. *Refer to "ntldr."*
- 8 C. The `ntbootdd.sys` file contains the SCSI driver for your SCSI adapter if the BIOS on the adapter has been disabled. *Examine "ntbootdd.sys."*
- 9 B. You should enter Safe Mode to correct problems when your computer won't boot normally. *See "Safe Mode."*
- 10 D. The file that makes up the user portion of the Registry in Windows XP is `ntuser.dat`. *Review "Registry corruption."*

Chapter 2: Managing Common Problems

Exam Objectives

- ✓ **Recognizing and resolving Windows-specific printing problems**
- ✓ **Identifying and correcting auto-restart and blue-screen errors and system lockup problems**
- ✓ **Knowing and fixing device driver failures**
- ✓ **Distinguishing and rectifying application install, start, and load failures**

In this chapter, you examine some of the common errors that you will encounter when servicing computers. I show you how to gather the information that you will need to properly diagnose and deal with the problem. You also look at several common errors and problems that you are likely to encounter when dealing with Windows-based systems.

These are some of the most common problems that you will run into at the OS level, so as a CompTIA A+ Certified Professional, this chapter will prepare you for these problems, and will help you on the exam.

Solving Windows-Specific Printing Problems

By now you should be familiar with the process of printing that was covered in Book III, Chapter 5. This section will look specifically at a few problems that are commonly encountered with printing from the OS perspective.

Even though people often say that we are moving toward a paperless society, people still do a substantial amount of printing every day. Because so many people are printing so often, everybody, at some point, has some sort of issue with printing. In the following sections, you examine stalled spoolers, wrong drivers, incorrect configuration parameters, and what you can do about them.

Dealing with a stalled print spooler

When working with a sewing machine, spooling allows you to take a pile of thread and put it on a spool for temporary storage. From the spool, the sewing machine could draw thread at a rate that was convenient for the machine, and would not require an operator to constantly deal with ensuring untangled thread was available for the machine. Print spooling does the same type of process for a printer.

Prior to print spooling being included with the OS, the application that was printing had to send each piece of data to the printer, and would not allow the user to do anything until the printing process was complete. The print spooler is used to quickly accept print data from the application, allowing control of the application to be returned to the user after a very brief spooling process. When the spooler takes the print data from the application, it stores the data in a temporary file, until the spooler has time to send the data out to the printer, or de-spool the data. From time to time, when dealing with Windows XP and other versions of Windows, you encounter a stalled print spooler. This seems to be a feature of the print spooler. After a period of printing, and printing well, the print spooler decides to stop responding to further commands. This has been a long-running issue for Microsoft that does seem to happen less and less with newer versions of the Windows OS.

You can identify a stalled spooler from the following:

- ◆ **Users are unable to add new jobs to the print queue.**
- ◆ **Nobody is able to remove jobs from the print queue.**
- ◆ **Existing jobs do not print.**
- ◆ **The print queue appears empty even though you have sent print jobs.**

If the print queue exhibits these symptoms, then you should be able to remedy them by restarting the print queue. This can be done by using either of the following methods:

- ◆ **Restart the spooler service by using the Services MMC snap-in.**

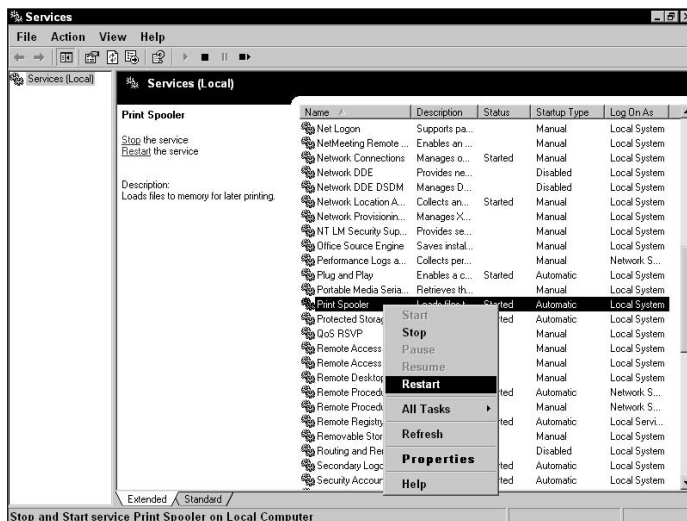
You can find a pre-saved Services console in your Administrative Tools folder. Choose My Computer⇨Control Panel⇨Administrative Tools⇨Services. Locate the spooler service in the list of services, right-click it, and choose Restart (as shown in Figure 2-1). You could also choose to Stop, wait for that action to complete, and then choose Start.

- ◆ **Restart the spooler service by using the `net` command.**

If your print spooler is suffering from any problem, restarting the spooler service is usually the correct answer.



Figure 2-1: Restarting services is easy from the Services Administrative tool.



If you want to use the command prompt, then you will be able to complete the same task at the server by opening a command prompt and typing

```
net stop spooler
net start spooler
```

or

```
net stop "print spooler"
net start "print spooler"
```



If you don't know the name of the service you want to start or stop, then type `net start`, and Windows XP lists the registered services.

Lab 2-1 provides practice restarting the print spooler via the command prompt. This lab requires a Windows 2000 Professional, Windows 2000 Server, or Windows XP computer. Lab 2-1 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

After the spooler has restarted, your problems should be gone, assuming you have configured the clients to use the correct driver.

Incorrect/incompatible driver for printing

Drivers are often one of the big problems for administrators. You notice that the printer is being heavily used, spitting out page after page of data, and you're feeling happy that this printing device is being put to such a good use. These feelings are then shattered when you examine one of the printed pages. You notice, up at the top of the page, a small band of text that looks

like it was typed over about 100 times. This should not greatly concern you, as the problem is not likely a major hardware problem with the printer, but rather one of the users on the network is using the wrong printer driver, and this is something you can fix easily.

The *printer driver* for your computer acts as a translation tool between your computer's OS and the printer. If you install the incorrect driver, then you have a poor translation taking place. The level of poor translation shows up in the printout. Sometimes an incorrect driver causes just some special characters to not print, and other times it causes a much larger problem, such as the one that was just described.

The solution to the problem is easy; just install the correct driver. That involves checking the model of the printer, getting the software for it, and installing the correct driver for the OS. If you don't have the driver disks or CD, you may be able to download the driver from the printer manufacturer's Web site or other Internet Web site.

Incorrect parameters

There are several parameters (settings) that can be configured for a printer that are specific to the individual printer. You can find these settings on the Device Options tab of the printer's Properties dialog box. Because the settings are specific to the device, I can't discuss them all in this book. What you see here are generic settings that apply to all printers. Of these, the ones that are most likely to cause problems are the port settings.

To get to a printer's Properties dialog box, choose Start⇨Printers and Faxes, and when the Printers and Faxes window opens, select the desired printer and right-click⇨Properties.

Port settings are found on the Ports tab. When choosing a port, you can pick from a local port like LPT1, a network port like \\<server>\<printer>, or a file to be redirected to a printer later. If you have specified an incorrect network location, then you will not be able to print, and if you print to a file, then you may not know where the file is. These can cause problems for people when they are printing, and you should verify that you have expected settings on this tab; for example, if you think you should be printing to LPT1, and you see that your selected port is COM1, then this is likely the cause of your printing problem.

Another problem that can be caused by settings on the Ports tab and the Configure Port button are the timeout settings. If these values are set too low, then you can have communication problems with your printer. The problem is caused by your computer not waiting long enough for your printer to answer any queries. If you are unsure of the correct port settings, then verify the port settings with information found in the printer's documentation.

Solving Boot Errors and Errors Requiring Restarting

With the number of computers that exist in most offices, and the way that the law of averages works, it seems like there is always a problem that requires or causes a reboot. In large offices, it can seem like the Welcome to Windows chime is always playing. In fact, so many support people have used rebooting as the first step in dealing with problems that some users now instinctively reboot even before calling in a support person. This is not always the best course of action, as it often removes information that can be used to locate the actual root cause of the problem.

Auto-restart errors

In many cases, auto-restart errors are caused by power-related issues, other hardware failures, or software configuration problems. If the computer auto-restarts, something is definitely wrong.

Some of the common issues that can cause auto-restarts include the following:

- ◆ **Service configuration:** When configuring services, you can choose an action to be taken when the service fails. One of the actions is to restart the computer. If you have a service that is failing and it is configured to restart the computer when it fails, then this would cause an auto-restart. This setting can be seen in Figure 2-2. More information on services can be found in Book VI, Chapter 3.

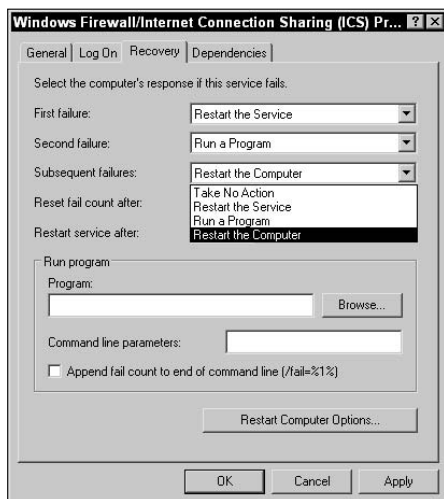


Figure 2-2: Auto-restarting a service after a failure is just one of the options.

- ◆ **CPU fan:** Most computers have a default thermal setting for the computer to cause a temporary shutdown. If the CPU fan fails, then the

CPU's temperature quickly rises to the point at which the motherboard interrupts the power to prevent damage to the CPU. Dust caked in and around the heat sink and fan has the same effect. A visual inspection of the CPU fan would let you know if it is running, or if there is a problem with dirt or debris. Book IV, Chapter 1, talks about ventilation methods and problems related to poor ventilation.

- ◆ **Power supply problem:** Power supplies have limited life spans, and when they fail it is often complete and immediate. In some cases, they can supply some power for a period of time, but as they heat up, they cause an interruption in the power, thus restarting the computer. It is easier to diagnose the problem when it is a complete failure of the power supply, rather than a power supply that fails when it hot; since to diagnose the latter, you need to have the power supply operating for a period of time, possibly days, before you see the problem. When in doubt, swap it out. Power supplies are not expensive and are easy to replace, and if it fixes the problem, you'll know it quickly. For more on power supplies, see Book II, Chapter 6.
- ◆ **Power source problem:** Some buildings are notorious for power spikes or drops. In some cases, the power may drop to a point that there is insufficient power for the computer, even though other electrical devices don't show any signs. If the problem happens regularly, then it will likely affect a number of computers.

There are a variety of tools for testing the quality of the power coming from the wall, and some low-end single-computer uninterruptible power supply (UPS) systems give you a display of the current voltage and provide line conditioning on a 120V line, down to 89V. The UPS can also supply some power when the voltage drops below that level. For more on power sources, see Book II, Chapter 6.

- ◆ **Bad memory:** When memory is the problem (the computer's, not yours), you may see the system running fine until the damaged or bad blocks of memory are used. At that time, the system spontaneously reboots itself or generates a Windows Stop error. There are testing tools to test for bad memory, so this situation should be easy to diagnose after letting the system run with the testing software working through the blocks of memory. As the size of the system memory goes up, so does the length of time that it takes to test it.
- ◆ **Network attack:** The TCP/IP network stack is very complex, and with services running over the stack, many software systems may have flaws. Many of the network attacks that may be performed on Windows systems cause various buffer overflows, which then cause the system to reboot. Worms like Zotob fall into this category. Tools like TDIMon and TCPView (www.sysinternals.com/NetworkingUtilities.html) can be used to view current TCP and UDP network activity on a computer, which can in turn be used to identify network attacks or problems; it is beneficial to learn how these tools work when you are not experiencing

a problem, since it will make it easier to identify non-standard traffic. In most cases, keeping your system up to date on security patches and virus definitions offers a high level of protection. The only sure protection from network attacks is to not be connected to the network. For more information about network security, read through Book IX.

- ◆ **Automatic Updates:** If your system is configured for Automatic Updates from Microsoft's Windows Update or Microsoft Update Web site, then you may find that your computer reboots while you are not there. If the computer is configured to install updates in the middle of the night, you may find when you return to it the next morning that it has rebooted. Usually, there are warning messages leading up to the reboot, but you might not even be there to see them. To solve this issue, check the Windows Event Logs and the settings on Automatic Updates in the Control Panel. As part of the new Windows XP Security Center, Windows suggests enabling Automatic Updates. Book IX, Chapter 3, has more information on the Windows Update process.

There are other reasons that unexpected reboots may happen, but these are the most common reasons that you might expect to see reboots happening on your system.



Automatic reboots of a system can have a variety of causes; make sure you are familiar with the list of the most common reasons.

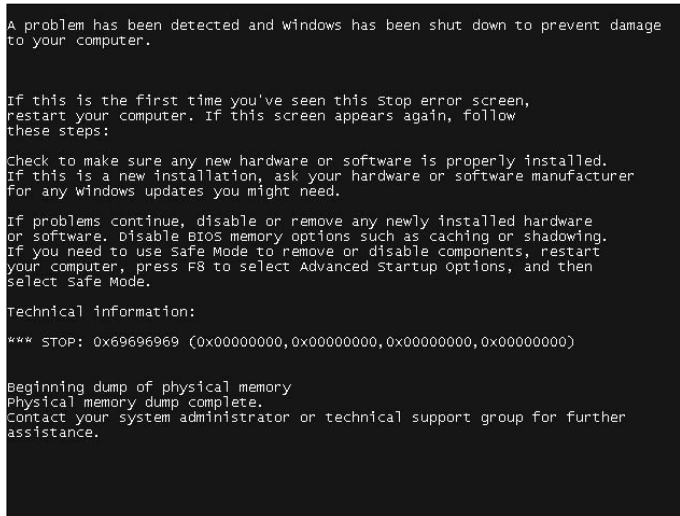
Blue-screen errors

Blue-screen errors (shown in Figure 2-3) are almost always related to driver, DLL, or configuration errors. *Blue-Screen errors*, or *BSoD (Blue Screen of Death) errors*, are officially *Stop errors*. In many cases, these will be related to a recent driver change prior to the last reboot. If this is the case, then you may be able to boot to the Last Known Good Configuration, which will restore the system configuration to the state it was in during the previous reboot. For more information about booting to the Last Known Good Configuration, see Book VII, Chapter 3.

The most common reasons for Stop errors are:

- ◆ **Service, application, or device errors**
- ◆ **Compatibility problems**
- ◆ **Hardware problems**
- ◆ **File system corruption or errors**
- ◆ **Compatibility issues with Firmware or BIOS**
- ◆ **Viruses**

Figure 2-3:
Stop errors are “affectionately” known as the Blue Screen of Death.



The image referenced in Figure 2-3 is a Stop error caused by an application, and the format of the text on the screen is not the same as the Stop error you may be more familiar with, which is discussed later in this section. Alternatively, the first two lines of the Stop error message tell you the name of one of the drivers that is involved, or at least the memory locations involved. Stop errors occur when one driver or application attempts to access the memory that is being used by another driver, and since drivers operate in the unprotected kernel memory space, this is very bad for system stability. When this happens, Windows does the only thing that it can think of to prevent further corruption — it stops everything. Because applications run in protected memory spaces, they will not be able to write to memory that is being used by another application, so Stop errors are only caused by access attempts in the kernel memory space.

The following text may appear on a standard Stop error:

```
*** STOP: 0x0000000A (0x802aa502, 0x00000002, 0x00000D00, 0xFA84001C)
IRQL_NOT_LESS_OR_EQUAL*** Address fa84001c has base fa840000 - i804prt.SYS

CPUID: GenuineIntel 5.2.c irql:1f      SYSVER 0xF0000565

Dll Base   Date Stamp - Name                               Dll Base   Date Stamp - Name
80100000   2be154c9   - ntoskrnl.exe                               80400000   2bc153b0   - hal.dll
80200000   2bd49628   - ncr710.sys                                  8025c000   2bd49688   - SCSIPTORT.SYS
...

Address  dword dump   Build [1381]                                     - Name
fe9cdae8 fa84003c fa84003c 00000000 00000000 80149905   - i8042prt.SYS
fe9cdaf8 8025dfe0 8025dfe0 ff8e6b8c 80129c2c ff8e6b94   - SCSIPTORT.SYS
fe9cdeb0 8013e53a 8013e53a ff8e6b94 ff8e6f60 ff8e6b94   - ntoskrnl.exe
...
```

Hexadecimal, the Base 16 number system, is used to display all of the error codes and memory addresses on the Stop error. The first line tells you the error type numerically (0x0000000A), then the memory address of the code that is making the access attempt (0x802aa502), the type of access attempted or error parameter (0x00000002, 0x00000D00), and the targeted memory address (0xFA84001C). This is followed by a line providing you with a text-based description of the error and, if possible, the name of the driver that was found at that location. This driver may have been the one that caused the error or the target of the error. After these lines is a section that lists the drivers that are found in the memory spaces near where the error occurred, then a section listing memory information for some of those memory spaces. If error was a result of a recent change to the system, undo that change or update your drivers.

Some common Stop errors you may encounter are:

- ◆ 0x0000000A IRQL_NOT_LESS_EQUAL
- ◆ 0x0000001E KMODE_EXCEPTION_NOT_HANDLED
- ◆ 0x00000024 NTFS_FILE_SYSTEM
- ◆ 0x0000002E DATA_BUS_ERROR
- ◆ 0x00000050 PAGE_FAULT_IN_NONPAGED_AREA
- ◆ 0x0000007B INACCESSIBLE_BOOT_DEVICE
- ◆ 0x0000007F UNEXPECTED_KERNEL_MODE_TRAP
- ◆ 0xC000021A STATUS_SYSTEM_PROCESS_TERMINATED

Make sure you are familiar with the basic type of information that is returned in Stop messages.



System lockup

System lockup is when your computer stops responding to any system functions or user input. If you are using Windows 2000 or Windows XP, system lockups should be rare. More likely, you'll have a period of slow responsiveness. If you are using older Windows operating systems, then you are likely to have complete system lockups. This has to do with how the operating systems work, and you may want to review Book VI, Chapter 2, to see how the Windows XP system architecture has been designed to prevent lockups.

If you are using Windows 2000 or Windows XP, then you may experience a runaway application or service that can cause the system to become unresponsive. This unresponsiveness may lead you to believe that your computer is locked up, when in fact it just doesn't have enough clock cycles to pay attention to you. If you press Ctrl+Alt+Del, you should be taken to the security dialog box (shown in Figure 2-4), where you will be able to launch Task Manager.

I have seen Windows XP lock completely up very few times, but there have been many other times when I have powered a machine off because it was slow to respond.

Figure 2-4:
The Security dialog box for Windows XP allows you to launch Task Manager.



The length of time that you wait for an unresponsive system is typically related to the number of critical unsaved documents that you had open when it became unresponsive. I have waited an hour to get the dialog boxes to close Microsoft Word and save my documents, prior to turning the computer's power off — turning the power off, rather than performing a graceful shutdown, is usually avoided due to the risk of disk corruption. This wait is not typical because in most cases, if the system is unresponsive to the point that Task Manager takes more than five minutes to open, then I usually hit the power button to kill the power to the computer.

Task Manager's Processes tab lets you see which application is the runaway. You just have to look for the application with the highest CPU value. This indicates which application is hogging the CPU. You may want to switch to the application and close it normally, but there is a good chance that this will not work. If you can't shut down the program normally, return to Task Manager, select the task, and click the End Task button. This should return your system to its previous level of responsiveness. In some cases, the application at fault will be a critical system service, like WinLogon, which cannot be stopped, so you will just have to reboot.

If you are using 16-bit Windows applications on Windows XP, then you are in for a different experience. Book VI, Chapter 2, examines how the system architecture lends to a lack of system stability. This is mainly due to OS components running in the same area of the system that 16-bit applications execute. When any 16-bit application crashes, it has the potential to lock the 16-bit OS components as well, which may lock up an entire NTVDM (NT Virtual DOS Machine). If this happens, your only choice will be to terminate the NTVDM that has the problem.

Resolving Device Driver Errors

A device driver's most common error has already been discussed — a system Stop error. Other than that type of error, device drivers may fail to load for a number of reasons:

- ◆ The version of the driver that you're using might not be compatible with the OS version that you're using.
- ◆ The driver could be misconfigured, and some settings require changing.
- ◆ The driver might not be compatible with some other driver or application that is running on the system.

In most cases, the startup error for the driver will be listed in the Event Viewer's error logs. You can open Event Viewer from Start⇨Control Panel⇨Administrative Tools⇨Event Viewer; then click on one of the three log files to view the error logs. After locating the error in Event Viewer, you may use other operating system tools, such as Device Manager, to correct the problem. Typically, the fix for these issues involves a visit to the vendor's Web site and downloading the latest version of the driver or reviewing their knowledge base for a fix to the compatibility problem. In rare cases, it may mean changing to a different version of the device, possibly from a different manufacturer.

Application Install, Start, and Load Failures

From time to time, you will have applications that will not start. A few reasons for this problem include corrupted or damaged shortcuts, damaged Program Information File (PIF) settings, or corrupted memory space. In the following sections, I will take a closer look at these three types of errors, as well as application installation errors.

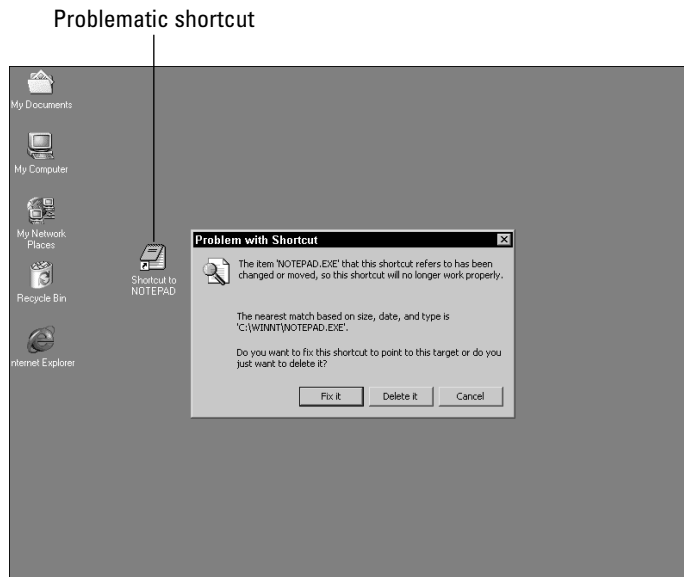
Corrupted shortcuts

When you create a shortcut to a program, the shortcut records information about the target file, such as size and creation date. If something happens to the original file, such as being moved to another directory, the shortcut will attempt to search the hard drive and repair the shortcut. If you are using Windows 2000, then you will be prompted to verify that the correct file was found, as shown in Figure 2-5. If you are using Windows XP, then automatic link tracking is enabled on shortcuts. What this means to you is that unless the file was deleted — in which case you will be asked to delete the shortcut — your shortcut links will always work, no matter how much you move a target file around.



Don't always choose to delete damaged shortcuts because Windows reports that a file is missing if your files are on a network server but you are disconnected from the network. If you reconnect to the network, your shortcut should work again. In the worst case, you have to re-create the shortcut by hand.

Figure 2-5: Shortcuts in Windows 2000 prompt you to confirm corrections for broken links.



Damaged PIF

In addition to improper locations, MS-DOS *PIFs* (*Program Information Files*) could also have inappropriate information in them. The PIF or shortcut to an MS-DOS application has information about the memory environment and reserved keystrokes. If this information has changed, then you will have to update the PIF. This can be done by right-clicking the PIF and choosing Properties. Which settings must be changed depends on your application, and Book VI, Chapter 2, deals with what those settings should be.

Corrupted memory space

Corrupted memory space occurs mostly with 16-bit Windows applications. Occasionally, 16-bit Windows applications can crash completely or partially when they are closed. When a 16-bit application crashes, it has a chance to corrupt the 16-bit Windows operating environment. If this happens, then you will find that you cannot open or launch 16-bit Windows applications, even though 32-bit Windows applications run fine.

If you're using Windows 9x, then you have to reboot to allow a new 16-bit Windows environment to load. If you are using Windows XP, then you can terminate the 16-bit Windows environment through Task Manager by following these steps:

1. **To open Task Manager, press Ctrl+Alt+Del and click the Task Manager button.**
2. **In Task Manager, click the Processes tab.**
3. **Locate the NTVDM that contains only `wow.exe`, which will be indented in the column, right-click it, and click End Task.**

The next time you load a 16-bit Windows application, it will reload the default 16-bit Windows environment.

Applications will not install

In some cases, you will find that applications just don't install. There are a number of reasons for this. In some cases, this is caused by the application performing a compatibility test, which your computer fails. If this is the case, then you should be notified of the compliance failure by the application. If your computer is less powerful than recommended, then it should be upgraded to support the application.

In some cases, specifically with Windows Installer applications, you may not be able to perform an installation if there is a previous installation that is still pending a reboot. The older setup applications may not install if they see another copy of `setup.exe` currently running. If this copy of setup is not expected to be running, then check Task Manager to see whether it is there. Rebooting the system will correct both of these issues.

In addition to compliance failure, there may be a problem with how the setup program was written. Very often, the setup program itself is a 16-bit Windows-based application. If you have recently had a 16-bit Windows application crash on your system, and you have not corrected the corruption of the 16-bit memory space by rebooting Windows 9x or terminating the NTVDM in Windows XP, then the setup program may fail.

Solving Other Problems

In addition to the problems already discussed in this chapter, a number of other things may go wrong. In the following sections, I examine some of those problems.

General protection faults

General Protection Faults (GPFs) are operating system–level errors. When applications are running, Windows prevents applications from interfering with each other by running them in their own memory space. However, some applications (and OS components) share a memory space — these are mostly 16-bit applications. When one of these components attempts to reference memory that does not belong to it, then Windows generates a general protection fault and attempts to prevent the improper reference by terminating the offending application.

Because this problem may affect some of the 16-bit OS components on the system, you should reboot after a GPF to ensure system stability. To reduce the occurrence of GPFs, you should reduce the number of applications that are running.

Illegal operation

The illegal operation error is similar to the GPF. In most cases, but not necessarily all, the illegal operation is a memory reference problem. When one 32-bit application attempts to reference an area of memory that belongs to another application or that it has somehow corrupted, then it generates an illegal operation. When this happens, Windows treats it as a rogue or damaged application and terminates it before it can cause damage outside its own memory space.

The big difference between illegal operations and GPFs is what components are in use and affected at the time. In most cases, you can recover from an illegal operation by re-launching the application. Reducing the number of open applications will reduce the chance of illegal operations because less space will be considered “out of bounds.”

Invalid working directory

When creating shortcuts of MS-DOS-based applications, you will configure a working directory. This directory is used to create any associated files that the application requires and to locate sub-applications that it requires to be able to perform its functions. If, after installing and configuring the application, this directory is removed or renamed, then you will see an error regarding the working directory (see Figure 2-6). As part of the error message, you are given a chance to launch the application. If you choose to launch the application, you should do so knowing that some functions may not work.

Figure 2-6:
If the working directory does not exist, you will be notified.



This error can be resolved by changing the working directory path in the PIF for the MS-DOS-based application. To do this, right-click the shortcut to the application (or the application itself) and choose Properties. In the Program tab, change the Working Directory setting to the correct directory, as shown in Figure 2-7. If the application requires support files, they may need to be restored. For more information about configuring PIFs see Book VI, Chapter 2.



Figure 2-7:
The Working Directory setting appears in all PIFs.

Optional device will not function

Many devices may be attached to your computer from time to time, but three devices that most computers have are a sound card, a modem, and a mouse. These devices all require an IRQ and an I/O address. When you are working with your computer, you may have other devices that are using resources that you require. If all of the devices on your computer are Plug

and Play, and you have a Plug and Play BIOS, then this will not be an issue. But if you are using legacy devices, you will have to configure IRQ and I/O settings for your devices. Sound cards usually use IRQ 5, and your COM ports (for your modem) are IRQ 3 and IRQ 4. If you are using a serial mouse, you want to make sure that it is not sharing its IRQ with COM 3 or COM 4. To find out more about configuring resources, read Book III, Chapter 4.

In addition to configuring the hardware support for the devices, you also have to load the appropriate driver for the device. This can be done through Add New Hardware in the Control Panel. For more information about installing devices on your computer, see Book VI, Chapter 1.

Terminate Stay Resident (TSR) programs and viruses

Microsoft DOS allows only one program to operate at a time, and this was often very limiting. *Terminate Stay Resident programs (TSRs)* are usually small useful programs that were designed to overcome the single application limitation of MS-DOS. When TSRs are launched, they are loaded into memory, and then they tell the OS that they are done, while remaining in memory, performing the task for which they were designed.

Book IX, Chapter 3, discusses viruses and some of the ways that they attempt to hide from detection. One of the ways that a virus spreads is by loading into memory as a TSR, and then they are able to start infecting other files on your system.

This is not to say the TSRs are to be avoided, but any program that runs without a user interface may be running when you are not aware. With Windows XP, you can see which applications are running in the background (as services) by bringing up the task list by pressing Ctrl+Alt+Del.

Your virus scanner has a list of known viruses that it will not permit to run. There is a chance that it will prevent others from running if they are doing things that are suspicious. If you are sure that the file is not a virus, then you can temporarily disable your virus protection. Many virus programs have an agent that runs all the time as a background service, looking for suspicious activity. In some cases, malware (which is not actually a virus, but still a program you don't want on your computer) will be allowed to run. Malware also runs in the background and should show up in your task list. Most malware applications or viruses try to disguise themselves so as to not be obvious about their presence. For more information about viruses, see Book IX, Chapter 3.

Getting an A+

This chapter goes over some problems that computer support personnel typically encounter. The problems include:

- ◆ Dealing with stalled print spooler by restarting services
- ◆ Fixing printing issues by replacing drivers or changing settings
- ◆ Identifying root causes for auto-restart and Stop errors
- ◆ Using Event Viewer to diagnose driver problems
- ◆ Dealing with application start problems in the form of shortcuts, PIF, or memory space issues

Prep Test

- 1** When the print queue is unresponsive, what should you do first?
 - A Reboot the server
 - B Reinstall the printer
 - C Restart the spooler service
 - D Redirect the printer to a remote print device

- 2** What command is used to restart a service from the command prompt?
 - A service.exe
 - B services.exe
 - C cmdrun.exe
 - D net.exe

- 3** You get a call from a user who is complaining that his document has bullets in it, but when he prints the document, he gets check boxes instead. What is the problem?
 - A He is using the incorrect printer driver.
 - B The application that he is using does not support bullets.
 - C His computer has a problem with the serial cable connecting his printer.
 - D He is using the wrong parallel cable.

- 4** Only 16-bit Windows applications will produce GPFs. True or False?
 - A True
 - B False

- 5** What program can be used to terminate programs that are running on your computer?
 - A Device Manager
 - B Computer Administration Console
 - C System Information
 - D Task Manager

6 Your computer has had a Stop error, and the first two lines read:

```
*** STOP: 0x0000000A (0x802aa502, 0x00000002, 0x00000D00, 0xFA84001C)
IRQL_NOT_LESS_OR_EQUAL*** Address fa84001c has base fa840000 - i804prt.SYS
```

Which of the following statements are true? (Choose all that apply.)

- A The type of error was IRQL_NOT_LESS_OR_EQUAL.
- B The error is due to the BIOS version.
- C The driver `i804prt.sys` was involved in the error.
- D The driver `i804prt.sys` is located at memory address `0x802aa502`.

7 What Windows XP function keeps your shortcut links from breaking?

- A LinkTracking
- B Shortcut Locator Service (SLS)
- C System Information
- D Task Manager

8 What tool should you check first when things are not working properly and you suspect an error, but you have not seen an error message?

- A System Information
- B WinMSD
- C Device Manager
- D Event Viewer

9 During your computer's boot process, your computer reboots automatically. There is no sign of a blue-screen error. What are the most likely issues? (Choose all that apply.)

- A Hardware problem with power supply
- B Driver issue
- C Service configuration
- D Memory error

Answers

- 1 C. Restarting the print spooler service should fix problems with the print queue. If it does not, then you may have to restart the server. *See "Dealing with a stalled print spooler."*
- 2 D. `net . exe` is the command that can stop and start services from the command prompt. *Review "Dealing with a stalled print spooler."*
- 3 A. He is probably using the incorrect printer driver. The problem should go away when the driver is updated. Most printers are attached by parallel cables. *Check out "Incorrect/incompatible driver for printing."*
- 4 B. This is false. 16-bit applications will GPF more often and are responsible for most GPFs, but not all of them. *Peruse "General protection faults."*
- 5 D. Task Manager can terminate applications or background processes that are running. Device Manager can be used to configure hardware devices. System Information can provide status and configuration information about most areas of your computer. I made the Computer Administration Console up. *Take a look at "System lockup."*
- 6 A, C. The error is of type 0x0000000A: IRQL_NOT_LESS_OR_EQUAL. The error involves `i804prt . sys`, which is found at memory address 0xFA840000. There is no information suggesting that the problem may involve the system BIOS. *Peek at "Blue-screen errors."*
- 7 A. Link tracking is the OS function that keeps track of the executables that shortcuts refer to. (SLS was a made-up term.) *Look over "Corrupted shortcuts."*
- 8 D. Event Viewer is one of the first utilities that you should use when trying to find out why your system is not responding correctly. WinMSD is the old term used for the System Information tool. Device Manager would list devices that did not start up correctly but would be unrelated to some errors that your system may experience. *Study "Resolving Device Driver Errors."*
- 9 A, C, D. The only item that it might not be is a driver issue. Most of the time, if there is a driver issue, the driver will either not start up or will trigger a Stop (blue-screen) error. *Refer to "Solving Boot Errors and Errors Requiring Restarting."*

Chapter 3: Preparing for Disasters with Disaster Recovery

Exam Objectives

- ✓ Managing emergency boot disks
- ✓ Booting Windows into special modes
- ✓ Understanding the Emergency Repair process
- ✓ Performing backups and restores

Even with the best planning, things go wrong. And when they do, you need to know your options for dealing with the situation. In this chapter, I show you the available options for recovery that are built into the OS. I also introduce you to the concept of a recovery partition and show you how to use a recovery CD.

Small issues that arise may be dealt with on the fly, if you know what tools are available to you. This chapter will review the tools that you will need to know. When a large issue appears, it is critical to have a backup, and know how to restore files from it. This chapter also will show you how to perform both of these processes.

Working with Boot Disks

If you find yourself with a computer that will not boot due to an operating system problem, then you will want to use a *startup disk* or a *boot disk* to get your system back into a bootable state. The following sections discuss boot disks for Windows 95, Windows 98, and Windows XP. A boot disk is a disk that has system files and is capable of starting the operating system on your computer.

Boot disks for Windows 9x

There are several different versions of the Windows 95 boot disk, one for each version of Windows 95: Retail/Upgrade, Plus Pack, OSR1, and OSR2. Keep in mind that when you are attempting to repair a system, some functions require that you use the correct version of the disk.

You can create a boot disk with just a couple of steps. This procedure is the same for both Windows 95 and Windows 98.

1. Open the Control Panel by choosing Start⇨Settings⇨Control Panel.
2. Double-click Add/Remove Programs.
3. Click the Startup Disk tab and then click the Create Disk button.

You will be prompted to provide a blank, unformatted disk, as shown in Figure 3-1.

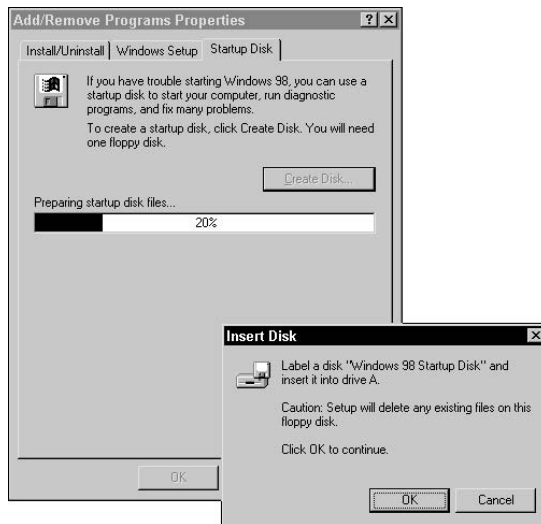


Figure 3-1:
Creating a
boot disk is
an easy
process.

4. Slide a fresh disk into the floppy disk drive and then click OK.

In a few minutes your disk will be ready. When it is done, the Startup disk tab will reset to display the Create Disk button.

5. Remove the disk and click OK to close the Add/Remove Programs dialog.

The basic boot disk for Windows 95 is the same across all versions of Windows 95. This disk contains the files listed in Table 3-1. These files represent the minimum number of files that Microsoft feels you need to recover or re-install a Windows 95 system.

Table 3-1

Common Windows 95 Boot Disk Files

<i>File</i>	<i>Description</i>
<code>fdisk.exe</code>	Used to modify the partition table on your drive. This is useful if you need to destroy and re-create the partitions on your disk.
<code>format.com</code>	Used to format existing partitions or new partitions.

<i>File</i>	<i>Description</i>
<code>edit.com</code>	Used to edit text files on your disk. These files include batch, configuration, and Registry import files.
<code>scandisk.exe</code>	Used to check the drive for bad clusters and other disk-related problems.
<code>regedit.exe</code>	Used to export or import Registry settings. This is the same file that is used graphically under Windows, but can be used here with a series of optional switches.
<code>sys.com</code>	Copies the required files needed to make a disk bootable. This can be used if the boot files on your hard drive have become damaged.
<code>attrib.exe</code>	Changes file attributes.
<code>chkdsk.exe</code>	Checks the disk for errors in directory tables. This was the precursor to <code>scandisk.exe</code> .
<code>command.com</code>	The 16-bit real-mode OS kernel.
<code>debug.exe</code>	A kernel level debugger for diagnosing low level errors. This is rarely used.
<code>drvspace.bin</code>	The code to handle disk compression.
<code>ebd.sys</code>	A flag file that lets the computer know that this is an operating system created boot disk. ebd stands for emergency boot disk.
<code>io.sys</code>	The real-mode boot loader.
<code>msdos.sys</code>	Real-mode boot settings.
<code>uninstal.exe</code>	The uninstaller for Windows 95. If a backup of system files was chosen during the installation of Windows 95, then this will allow you to uninstall Windows 95. This file is not on OEM versions of Windows 95 because they were not supposed to be installed over another operating system.
<code>scandisk.ini</code>	Preferences for Scandisk.

The files that truly make these versions of the boot disk different are `io.sys`, `command.com`, and `drvspace.bin`. Each of these files is version-dependent. If you are using one version of the boot loader, then you must use a matching version of `command.com`. If you do not, you get an error message telling you that `command.com` is the wrong version and your system will not boot. Because some system tools are updated when the Windows 95 Plus Pack is installed, it has a boot disk. Boot disk versions are

- ◆ Windows 95 Retail/Upgrade
- ◆ Windows 95 Plus Pack
- ◆ Windows 95 OSR1
- ◆ Windows 95 OSR2

If you are attempting to resolve boot or missing `command.com` problems and you don't know the exact version of Windows 95 that is being used, then you may have to try all versions of the boot disk to resolve the problem.

Windows 98 has two versions of the boot disk: One for Windows 98 and one for Windows 98 Second Edition (SE). These disks are primarily the same with the exception of the version of `command.com` that is used. Both Windows 98 boot disks start by creating a large RAM disk and extracting a compressed file (`ebd.cab`) to the RAM disk. A RAM drive is a portion of RAM that is treated like a hard drive and assigned a drive letter. Key files that are different from the Windows 95 boot disk are summarized in Table 3-2.

Table 3-2 Windows 98 Boot Disk Files	
<i>File</i>	<i>Description</i>
<code>Aspi2dos.sys</code>	
<code>aspi4dos.sys</code>	
<code>spi8dos.sys</code>	
<code>aspi8u2.sys</code>	
<code>aspicd.sys</code>	
<code>btcdrom.sys</code>	
<code>btosm.sys</code>	
<code>flashpt.sys</code>	
<code>oakcdrom.sys</code>	Several CD-ROM drivers for the Windows 98 Command Interpreter. Windows 98 actually loads all of the drivers in an attempt to find one that works.
<code>autoexec.bat</code>	A boot file that loads drivers after the <code>config.sys</code> file is processed.
<code>command.com</code>	The 16-bit real-mode OS kernel.
<code>config.sys</code>	A file that loads initial device drivers and is used to build the boot menu that loads when booting from the disk.
<code>drvspace.bin</code>	This code handles disk compression.
<code>Ebd.cab</code>	This is a 266K compressed file that contains most of the troubleshooting and diagnostic tools that are found on the boot disk, such as <code>attrib.exe</code> , <code>chkdsk.exe</code> , <code>debug.exe</code> , <code>edit.com</code> , <code>ext.exe</code> , <code>extract.exe</code> , <code>format.com</code> , <code>help.bat</code> , <code>mscdex.exe</code> , <code>restart.com</code> , <code>scandisk.exe</code> , <code>scandisk.ini</code> , and <code>sys.com</code> . This file is decompressed onto a RAM drive that is created during the boot disk boot process.
<code>extract.exe</code>	Used to decompress the <code>ebd.cab</code> file.

<i>File</i>	<i>Description</i>
<code>fdisk.exe</code>	Used to modify the partition table on your drive. This is useful if you need to destroy and re-create the partitions on your disk. This version can work with FAT32 partitions and is aware of NTFS partitions.
<code>findramd.exe</code>	Responsible for finding the drive letter that is assigned to the RAM drive that is created by <code>ramdrive.sys</code> .
<code>himem.sys</code>	The DOS-mode high memory manager.
<code>ramdrive.sys</code>	The driver file for the RAM drive.
<code>setramd.bat</code>	Searches for the first drive letter that may be used for the RAM drive.

One of the other main differences between the Windows 95 and Windows 98 boot disks relates to CD-ROM support. Windows 98 attempts to load several different CD-ROM drivers while attempting to find one that works in your computer. This makes the process of loading a CD-ROM driver easier than with Windows 95, which required you to locate and manually load the CD-ROM driver.

Since the Windows 98 boot disk is very good at locating and loading CD-ROM drivers, many people are now using it as a generic disk to load support for a local CD-ROM even if they are not planning on installing Windows 98. The boot disks for Windows ME are similar to those for Windows 98.

Windows XP

Windows XP doesn't truly have a boot disk, but a disk can be formatted with Windows XP, and the Windows XP boot files can be copied over to the disk. This disk can then take the place of a corrupted boot sector on a Windows XP system. Book V, Chapter 6, describes the files that are required for booting Windows XP and provides details about each. Book VII, Chapter 2, looks at boot issues related to these files and how to resolve the boot issues. The files required on this disk are

- ◆ NTLDR
- ◆ NTDETECT.COM
- ◆ BOOT.INI
- ◆ NTBOOTDD.SYS (if it exists)

The Windows XP boot disk is only used as a replacement for a damaged or corrupted hard drive boot sector.



Using F8 Options during Boot-up

With Windows 9x, you can press the F8 key during boot-up when you see the words “Starting Windows 9x” on the screen. On some computers, Windows 98 requires that you press and hold the Ctrl key immediately after the POST process completes rather than using the F8 key. With Windows 2000 or Windows XP, a message will appear at the bottom of the splash screen at the start of the boot process that tells you to press the F8 key to see the advanced boot options. Any of these methods will bring you to a boot menu, where you will be able to choose Safe Mode or Safe Mode with Networking. Windows 2000 and Windows XP also has a Safe Mode with Command Prompt option that loads `cmd.exe` as your shell application.

When you access the F8 boot menu at system startup, you reveal several advanced boot options depending on your OS:

- ◆ Safe Mode
- ◆ Safe Mode with Networking
- ◆ Safe Mode with Command Prompt
- ◆ Enable Boot Logging
- ◆ Enable VGA Mode
- ◆ Last Known Good Configuration
- ◆ Directory Services Restore Mode
- ◆ Debugging Mode
- ◆ Disable Automatic Restart on System Failure
- ◆ Step-by-Step with Confirmation
- ◆ Command Prompt
- ◆ Previous version of MS-DOS

Booting into Safe Mode

If the computer hangs or crashes and reboots during the boot process, then it may be a device-related error. If either a Windows 9x, Windows 2000, or Windows XP computer fails to boot properly, and you think that the problem is related to a service or driver that is loading, then you might be able to boot the computer into Safe Mode.

Windows NT 4.0 does not have a Safe Mode boot, but it does have Enable VGA (Video Graphics Array) mode that you can select from the main boot menu. This is similar to Windows XP Enable VGA Mode which loads all system drivers, with the exception of the video driver. The idea of this option is that with the exception of having the video settings set wrong, the



computer will boot, and report that some drivers had failed to load, but it will boot.

When Windows 9x detects that it has failed to boot properly, it automatically attempts to boot into Safe Mode during the next boot.

When booting into Safe Mode, the operating system skips `config.sys` and `autoexec.bat`, as well as any drivers that are not considered crucial to the boot process. This means that you get a VGA driver providing video in a daring 640 x 480 resolution (16 colors), a keyboard, and a mouse. Initially, you will be presented with a warning telling you that you are entering Safe Mode, after which, the words *Safe Mode* appear in each corner of your screen, as shown in Figure 3-2. Outside of these devices, you will have very few working devices — for example, no sound cards, scanners, or CD burners. With just the basics running, you should be able to identify the driver that did not load properly and correct the problem.




Figure 3-2: Safe Mode should be easy to spot, with the wording in each corner.

After entering Safe Mode, the next three times you reboot the system, Windows automatically creates a `bootlog.txt` file on the root of your drive. If your next attempt at a normal boot doesn't succeed, you can read the `bootlog.txt` file and try to find the misbehaving driver or service there.

Windows 2000 and Windows XP also support a Safe Mode with Command Prompt. In Windows 2000 and Windows XP, this mode boots the Windows GUI (graphical user interface) but opens `cmd.exe` as your shell. This allows you to perform many actions to correct your problem (by using utilities like `chkdsk.exe` or the Registry Editor), and when you are done, exiting the Command Prompt reboots your system.

Both Windows 2000 and Windows XP have a boot option for Safe Mode with Networking. This mode is useful when you need to download new drivers for a failed device, or need to access files that may be stored on your server. With the exception of the loading of the networking drivers, this is a standard Safe Mode boot.

Tracking the boot with a bootlog

The `bootlog.txt` file is very useful for troubleshooting boot problems. It logs all drivers, services, and resources as they are loading. This file can then be referenced to locate problems with the boot, especially since the system does not give you much indication of what is happening during the boot. If a `bootlog.txt` file already exists then the existing (previous) one is renamed to `bootlog.prv`.

The structure of a `bootlog.txt` file from Windows 98 is shown in Listing 3-1.

Listing 3-1: A Windows 98 bootlog.txt File

```
[00055731] Loading Device = C:\WINDOWS\COMMAND\DRVSPACE.SYS
[00055731] LoadFailed = C:\WINDOWS\COMMAND\DRVSPACE.SYS
[00055731] Loading Device = C:\WINDOWS\HIMEM.SYS
[00055732] LoadSuccess = C:\WINDOWS\HIMEM.SYS
<Lines Deleted>
[0005578C] SYSCRITINIT = VMM
[0005578C] SYSCRITINITSUCCESS = VMM
[0005578C] SYSCRITINIT = MTRR
[0005578C] SYSCRITINITSUCCESS = MTRR
[0005578C] SYSCRITINIT = VCACHE
[0005578C] SYSCRITINITSUCCESS = VCACHE
[0005578C] SYSCRITINIT = DFS
[0005578C] SYSCRITINITSUCCESS = DFS
<Lines Deleted>
[0005578D] DEVICEINIT = VMM
[0005578D] DEVICEINITSUCCESS = VMM
[0005578D] DEVICEINIT = MTRR
[0005578D] DEVICEINITSUCCESS = MTRR
[0005578D] DEVICEINIT = VCACHE
[0005578D] DEVICEINITSUCCESS = VCACHE
[0005578E] DEVICEINIT = DFS
[0005578E] DEVICEINITSUCCESS = DFS
<Lines Deleted>
[0005578D] Dynamic load device bios.vxd
[0005578D] Dynamic init device BIOS
[0005578D] Dynamic init success BIOS
[0005578D] Dynamic load success bios.vxd
<Lines Deleted>
[0005579F] Dynamic load device C:\WINDOWS\system\IOSUBSYS\apix.vxd
[0005579F] Dynamic load success C:\WINDOWS\system\IOSUBSYS\apix.vxd
[0005579F] Dynamic load device C:\WINDOWS\system\IOSUBSYS\cdfx.vxd
[0005579F] Dynamic load success C:\WINDOWS\system\IOSUBSYS\cdfx.vxd
<Lines Deleted>
Initializing KERNEL
LoadStart = system.driv
```

```
LoadSuccess = system.drv
LoadStart = keyboard.drv
LoadSuccess = keyboard.drv
LoadStart = mouse.drv
LoadSuccess = mouse.drv
<Lines Deleted - This portion starts a shutdown.>
Terminate = User
Terminate = Query Drivers
EndTerminate = Query Drivers
Terminate = Unload Network
EndTerminate = Unload Network
Terminate = Reset Display
EndTerminate = Reset Display
EndTerminate = User
```

The `bootlog.txt` file in Listing 3-1 lists the devices that were loaded and the order in which they were loaded. In this case, `drvspace.sys` failed to load while `himem.sys` loaded. They were followed by Virtual Memory Manager, the Disk Cache, and the Distributed File System driver. Look through the listing above and see what other drivers or devices you can identify. When checking the file, you should look for items that list `LoadFailed`, since it will be failures that are related to current problems.

When choosing the Enable Boot Logging option from the Windows XP F8 menu, you will also create a log file called `ntbtlog.txt`, which is saved to the `winnt` or `windows` directory. Unlike the Windows 9x `bootlog.txt` file, this file is appended to when performing logged boots, but has a similar structure to the Windows 9x `bootlog.txt` file.

Using the Last Known Good Configuration

If you have just updated a driver on your Windows NT-based computer and your computer will not boot, you can try booting into Safe Mode to troubleshoot the issue, but if the driver update is causing problems with Safe Mode as well, you still have options. Later in this chapter, you will learn about the Recovery Console or the Emergency Repair Process, but in some cases it is quicker to restore the Last Known Good Configuration.

There is a registry key that stores the load or startup settings for all devices and services on your computer, including the path to the driver that is to be used. The registry key is `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet`. By choosing to restore the Last Known Good Configuration from the F8 boot menu you can quickly and automatically restore this section of the registry with a copy that was used during your last successful boot. If your computer had been running for a long time with many potential changes to this registry key, those other changes will be lost while you are repairing the boot problem. You will have to decide if the loss of those other changes is worth the quick restoration of a bootable computer.

Other Windows 2000 and Windows XP boot options

There are a few other advanced boot options for Windows XP which you will rarely use, so they will be briefly covered here. These options are

- ◆ **Directory Services Restore Mode (for Windows 2000 or Windows 2003 Domain Controllers):** This boot option will never be used with Windows XP, even though it is in the options menu, as it is designed to restore the Active Directory database which would be running on a server
- ◆ **Debugging Mode:** This option is for use with external debuggers, which you may be asked to set up during troubleshooting sessions with Microsoft Professional Services after placing a support call
- ◆ **Disable Automatic Restart on System Failure:** Prevents drivers that have been configured to initiate a reboot when their load fails, from causing the system to reboot. This option would be used to allow the system to complete its boot, so that you could perform troubleshooting from within the OS

Entering MS-DOS Mode

Windows 9x supports a non-graphical boot. In the boot menu described in the section “Using F8 Options during Boot-up,” DOS Mode is accessed by choosing Command Prompt or Safe Mode Command Prompt. If you choose Command Prompt, Windows continues to boot normally but stops just before it would normally launch `win.com`. In Safe Mode Command Prompt, `config.sys` and `autoexec.bat` are not processed, and the system doesn’t load the `himem.sys` memory manager. This boot stops prior to launching `win.com` as well.

The Command Prompt boot follows the normal Windows 9x boot process, but it stops prior to loading the GUI. From the Command Prompt, you can copy files and import or export Registry settings. You can also launch Windows with or without debugging switches by running `win.com`.

If your system boots up when using Safe Mode or Safe Mode Command Prompt, but will not boot when using a normal boot or with Command Prompt, then the problem is related to a driver that is loaded in `config.sys`, `autoexec.bat`, or through the registry. Windows 9x has another boot option named Step-by-Step with Confirmation. As Windows 9x processes each line of `config.sys` and `autoexec.bat`, it will prompt you to see if you want to load or skip the line. The last item that it will ask about is `win.com` to launch the Windows 9x GUI.

Emergency Repair

When things go wrong with the boot of Windows NT, Windows 2000, or Windows XP, you can try booting into Safe Mode to repair the error, but there are times when the system will not boot at all. At these times, you can make a final attempt to resurrect the system or start the installation. If you choose to attempt the Emergency Repair Process, you should be prepared by creating an Emergency Repair Disk prior to your computer experiencing problems, and you should be familiar with the repair process.

Emergency Repair Disk (ERD)

Windows 2000 and Windows NT both allow you to create an *Emergency Repair Disk (ERD)*. The ERD contains a small replacement version of the system account database and a major portion of the system Registry. This disk can restore the settings and list of the devices or services on the system. It can also replace the account database with a copy of the account database that existed at the time that you created the ERD.



The ERD is not required when performing an Emergency Repair with Windows XP and newer Microsoft operating systems.

With Windows NT, the Emergency Repair Disk is created by the `rdisk.exe` utility. You can run `rdisk.exe` from a command window or the `run` command. This utility updates several system files and places copies of them in a special directory, `%winroot%\repair`, which can be copied to a floppy disk after updating. The normal `rdisk` command does not update the `SAM` and `SECURITY` files (see Book VI, Chapter 4), but these can be updated by using `rdisk /s`.

If you're using Windows 2000, the Emergency Repair Disk is created with the `backup` utility found at Start⇒Accessories⇒System Tools⇒Backup. When you open this utility, there is an option to create an Emergency Repair Disk on the Welcome tab.



With each new version of the Windows OS, the registry has become larger. Until now it was not able to fit onto a floppy disk. Windows XP does not use an ERD because it cannot fit the required files onto a disk. The Windows XP Emergency Repair Process stores its backup copy of the registry files in the Windows directory, so unless you have hard drive failure, these files will be available when they are needed.

Emergency Repair Process

The Emergency Repair Process begins with booting the system with the Windows OS CD and starting the setup process. After starting the setup process, the first question you are asked is if you would like to perform an

installation of Windows 2000 or Windows XP, or if you would like to perform an Emergency Repair.

After choosing to perform an Emergency Repair, you are asked whether you would like to launch the Recovery Console (discussed in the next section) or perform an Emergency Repair. If you choose to perform an Emergency Repair, you are asked whether you have an Emergency Repair Disk. If you have your ERD, then Windows will use the files on the disk; but if you do not have your ERD, the Emergency Repair Process searches your drive for a Windows 2000 or Windows XP installation. If it finds one, it uses the repair information found in the `%winroot%\system32\repair` directory. You then have four choices for the Emergency Repair Process:

- ◆ **Inspect boot files**
- ◆ **Inspect startup environment**
- ◆ **Inspect system files**
- ◆ **Inspect Registry**

If you choose to inspect the boot files or the startup environment, the Emergency Repair Process will automatically fix problems that may exist with either the boot sector or start files, such as `boot.ini`.

If you choose to inspect the system files, the Emergency Repair Process will compare every file that makes up the Windows installation to the originals that are found on the Windows CD. If it finds any files that are different, the files may be corrupted, so the Emergency Repair Process will ask you if you want to replace the file with the original copy from the Windows CD.

If you choose to inspect the Registry, you will then have three choices:

- ◆ **Load the SYSTEM key**
- ◆ **Load the SOFTWARE key**
- ◆ **Load the SAM and SECURITY keys**

The Inspect Registry option doesn't so much inspect the sections of the Registry as replace them. The `SYSTEM` and `SOFTWARE` sections of the Registry contain settings for services, devices, and applications that are operating on the system. Most services store their settings in the `SYSTEM` key, but some (for Windows logo requirements) are also stored in the `SOFTWARE` key. The `SOFTWARE` key is used primarily for applications.

The `SAM` and `SECURITY` keys are replaced as a set. They contain the entire local account database. Replacing the account database on a member server will affect only that one computer. If you make the replacement on your Windows NT primary domain controller, then you will have replaced your

entire domain account database. On Windows 2000 domain controllers, the local account database is not used.

Rather than choosing to perform the emergency repair process, you could choose to launch the Recovery Console, which contains its own repair and troubleshooting tools. For more on the Recovery Console, proceed to the next section.

Recovery Console

The Recovery Console is a command-line base OS that can be used to help diagnose and repair problems with your installation of Windows 2000 or Windows XP. This command-line interface comes with several commands that look familiar to MS-DOS or Windows commands, as well as some that are unique to the Recovery Console. Even the commands that are named the same as MS-DOS commands tend to perform different functions.



Keep in mind that this is not MS-DOS, nor is it Windows XP. It is the Recovery Console, a unique OS used to repair Windows XP.

You can get to the Recovery Console through the emergency repair process, or you can install it into your boot menu by re-running the Windows XP setup program with the `cmdcons` switch:

```
<winnt_src>\i386\winnt32.exe /cmdcons
```

To initially get into the Recovery Console, you have to log into it. The Recovery Console checks for copies of Windows XP on your system and lets you log in to one that you choose. In order to log in, you need the name and password of the local Windows XP administrator account. After providing that information, you will have access only to the files that make up that copy of Windows XP.

You have access only to the `winnt` or `windows` directory. By limiting your access, Microsoft increases security and restricts the amount of damage/harm you can do. If you are booting to the Recovery Console, then your computer is likely having major problems that you were not able to solve through other means.

The Recovery Console offers you several file management tools. Of these, the most notable are `copy`, `del`, and `ren`. All three of these commands work with only one file at a time and do not support wildcards like their Windows XP equivalents. `format` is the other major file management command at the console. `format` enables you to specify the file system for your partition — FAT, FAT32, or NTFS.

Several commands let you manage the disk and boot sectors on your drive: `chkdsk` for checking your disk for errors; `diskpart` for repartitioning, much

like `fdisk` would; and `fixboot` and `fixmbr` for rewriting either the boot sector or the Master Boot Record to resolve system startup errors.

In addition to using these commands, you often need to control services and devices during the troubleshooting process. To help with this, you can use the following commands:

- ◆ `listsvc` lists the available services and devices
- ◆ `disable` disables selected services and devices at boot
- ◆ `enable` enables selected devices and services at boot

A couple of unclassified commands include the `logon` and `systemroot` commands. `logon` allows you to log on to another local copy of Windows 2000 for systems that dual boot. `systemroot` switches you to the Windows directory that you are logged onto.

With this assortment of tools, you should be able to get an unbootable copy of Windows XP bootable again. Don't expect to perform much troubleshooting through the Recovery Console, but if your server is having STOP errors on boot, the Recovery Console will allow you to get around them. When you are once again running a full copy of Windows XP, you should be able to complete your troubleshooting procedures.



Lab 3-1 gives you some practice working with the Recovery Console. For this lab, you need a computer running Windows 2000 Professional or a version of Windows 2000 Server. In addition, you need an installation CD for the operating system and four blank floppy disks. Even though your computer may be able to boot directly from the CD-ROM, this lab makes installation boot disks for Windows 2000 and uses them to initiate the setup routine. If you choose to boot from the CD-ROM, you may have to change the settings in your CMOS, and you can skip down to Step 5 in this lab exercise. Lab 3-1 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Understanding Backup and Restore

Ensuring data recoverability involves performing a *backup* (normally to another type of medium) in addition to storing the file on the hard drive. This medium could be (for example) a floppy disk, another hard drive, a tape drive, a Zip drive, or a CD-ROM.

In the event that you have lost a single file or lost all of the files on your hard drive, you will want to perform a restore. You will need to have your backup files or tapes available and will use the Windows XP backup utility to restore your files.

Backing up your computer

The Windows 2000 and Windows XP backup software is greatly improved over Windows NT 4.0 backup software. The new backup software picks up where its predecessor left off, allowing you to schedule regular automatic backups and to back up to nearly any medium you like (including, but not limited to, tape drives). Even though the tool is new, Microsoft has kept the name of `ntbackup.exe`, and supports all of the old command line switches for backward compatibility.

The backup utility lets you create a backup of all of the local files on your disk drives, as well as the system state. Open files, such as the Registry and system databases, cannot normally be backed up on your computer, but system state allows you to back up these files by using new file locking methods.

Each organization decides how often it needs to back up each of its computers. This decision is often based on the size and ease of use of the backup media and the value that the organization puts on its files. For example, the loss of a week's worth of invoice records may represent a large amount of revenue to a company, or the loss of a day's worth of rental records at a video store may put a substantial portion of the video store's inventory in jeopardy, since they will not have a record of where their movies are. These types of costs are weighed against of the cost of decent backup hardware. Some organizations may back up important data once a week, once a day, or several times a day, depending on the perceived cost of data loss.

You start a typical backup by launching the Backup utility, which can be done in one of two ways:

- ◆ **By running `ntbackup.exe`.**

All of the options for the backup utility are available via the command line interface. To get a complete list of the options and their command line switches, type `ntbackup.exe /?` at a command prompt.

- ◆ **By choosing `Start` → `All Programs` → `Accessories` → `System Tools` → `Backup`.**

Either of these methods launches the backup utility. The first time you launch it, it launches in Wizard mode, which can be used for backup, and will allow you to back up files using simplified or default settings. This section will guide you through the process using Advanced mode, which allows you set detailed options for your file backup or restore.

If you don't want to launch it in Wizard mode, then uncheck the Always start in wizard mode check box in the Backup or Restore Wizard dialog box and click the Cancel button or the Advanced Mode hyperlink. The next time you launch the backup utility, it will launch in Advanced mode instead of in Wizard mode. When the utility is in Advanced mode, you can run the Backup, Restore, or ASR Wizards. You can also choose the appropriate tabs



to perform a manual backup or restore. The last tab available in the program is the Schedule Jobs tab, which is used to schedule automated backup jobs.

To schedule a basic backup of your system without using the Wizard mode, follow these steps:

1. Choose Start⇨All Programs⇨Accessories⇨System Tools⇨Backup.

The Backup Utility dialog box should open.

2. If the Backup or Restore Wizard opens, deselect the Always start in wizard mode check box and then click the Advanced Mode link.

3. Click the Backup tab.

4. Select all of the drives, folders, or files that you want to back up, as well as the System State, by clicking in the Selection boxes to the left of the names.

Backing up the System State will back up core operating system files, including files that are open. For Windows 2000 and newer Windows operating systems this option will include

- COM+ Class Registration database
- Boot files, including the system files
- Registry

If you are backing up a server, the System State may include some or all of the following, depending on what services are installed on the server

- Certificate Services database
- Active Directory database
- SYSVOL directory
- Cluster service information
- IIS Metadirectory

System State backups are required to fully restore a computer to its original state and identity on the network. If you are only concerned with the data and file system permissions, you do not need to back up the System State.

5. Select a location to back up to and then click the Start Backup button.

You can back up the files to another local drive or to a tape drive if there is one attached to the computer.

If you are saving to a file on a drive, then the default extension for a backup file is .bkf.

6. In the Backup Job Information window, click the Advanced button to set the advanced options for the job.



The main items to be aware of are

- *Verify Data after Backup:* This option rewinds the tape or rereads the backup file and compares its contents to each file that was backed up. You can disable this option if you find that some files managed by third-party applications become corrupted during the verification process. This is not normally the case.
- *Disable Volume Shadow Copy:* Volume Shadow Copy is a service that takes a snapshot of all open files on your drive to allow them to be backed up. There may be some third party applications that may have problems with the Volume Shadow Copy process, so you have the option of disabling this feature, but disabling the Volume Shadow Copy process will prevent some open files from being backed up, as the older file lock mechanisms will be attempted.
- *Backup Type:* You have five options, which are summarized in Table 3-3.

7. Once the Advanced options are set, click the OK button.

The Backup Job Information window reappears.

8. In the Backup Job Information window, click the Schedule button.

You are prompted to save your backup selections into a `.bks` (backup selections) file.

9. Select a location, name the file, and click Save.

This file can be stored anywhere on your drive, but for organizational purposes, I suggest that you save it in the same location as the `.bkf` file. Saving the backup selections file allows the utility to know what files are to be copied during each scheduled backup time. The `.bks` file can be loaded into the backup utility at any time you want the same selections for a backup.

After you click the Save button, the Task Scheduler service prompts you to enter your credentials. The default Windows Task Scheduler service is used when creating the scheduled automatic backup jobs that will appear on the Schedule Jobs tab of the Backup Utility.

10. Provide the credentials of a user who had the operating system right to perform a backup of the system. This would include members of the Administrators or Backup Operators groups. The system will use the credentials you supply when it automatically starts the scheduled backup process. After providing the credentials, click OK.

You are presented with the standard options for the Scheduling Service.



Operating system rights are similar to file system permissions, but they grant a username the ability to perform an action that is not directly related to an object, such as a file. You assign rights using the Local Security Policy or an Active Directory Group Policy Object (GPO). To learn more about rights, read through Book IX, Chapter 2.

11. Select when you would like the backup to run and then click OK.

Table 3-3 Backup Types	
<i>Name</i>	<i>Description</i>
Normal	Backs up all selected files and clears the Archive attribute so that the files can be selected by incremental or differential backups if they are modified. If time and storage capacity permit, normal backups are usually the most desirable.
Copy	Backs up all selected files but does not clear or modify the Archive attribute of the file. This allows you to perform a full backup of your files with the intention of giving the backup to another group, such as the finance department at month's end. Because the Archive attribute isn't touched, your other incremental or differential backups will still be valid and will work as normal.
Incremental	Backs up any selected files that have their Archive attribute set and then clears the Archive attribute. This means that only files that have changed since the last full or incremental backup will be backed up. In order to restore the files, you need the last normal backup and all incremental backups that have been taken since the normal backup. Even though you need several backups to perform a restore, each backup will be small in comparison to the normal backup.
Differential	Is similar to incremental backups in that it reduces the time it takes to perform the backup, it reduces the space required for the backup, and it relies on an existing normal backup. The difference with this backup is that the Archive attributes of files are not touched. So with each backup, more and more files are backed up. The benefit of this is that only the most recent normal backup and the most recent differential backup are required to perform the restore.
Daily	Backs up only the files that were changed today and does not touch the Archive attribute. If my normally scheduled backup runs at midnight, and at 4 p.m., I want to run a utility on my drive that might corrupt data, I can run a Daily backup. Because it backs up only the files that were modified during the day, it will be a quick backup; and because it doesn't touch the archive attributes, if the backup isn't needed, I can continue to use my normal routine for backing up and restoring files. It's like the backup didn't happen — much like the Copy type listed above.

By default, backups create a summary log that lists exceptions or the files that were not backed up. After scheduling and running a backup, you can log onto the computer by logging on as the backup user account that was specified when scheduling the job (Step 10 in the previous instructions). You can then open the backup utility and choose Tools⇒Report. This allows you to read the last 10 backup logs.

Restoring files from a backup

Now that you have seen the steps required to perform a basic backup of files on your system, you should become familiar with the steps required to restore those files. To restore files on your system without using the Wizard mode in the Backup Utility, follow these steps:

1. Choose Start⇒All Programs⇒Accessories⇒System Tools⇒Backup.

The Backup Utility dialog box should open.

2. If the Backup or Restore Wizard opens, deselect the Always start in wizard mode check box and then click the Advanced Mode link.

3. Click the Restore and Manage Media tab.

This tab has two panes. The left pane is used to select and catalog backup files you want to restore from and to display the directory tree in a backup file, while the right pane allows you view information about backup files that have been cataloged and to select items that you want to restore. This dialog should already have a backup file cataloged, which will be the one that you created in Step 6 of the backup process.

If you need to import or catalog a backup file that was created on another computer, or that you have manually deleted from this window, then right-click on the word File in the left-hand pane and select Catalog file, locate and select the backup file that you want to work with, and then click OK.

4. Double-click on the Backup Identification Label of the backup file that you want to work with.

5. Using the navigation controls in both panes, navigate through the list of files and folders that are found in the backup and select the files that you want to restore.

6. Select one of the three options for the Restore files to drop down menu.

- *Original location:* Restore folders and files to the same locations from which they were backed up.
- *Alternate location:* Restore folders and files to an alternate location, creating a duplicate of the directory structure that was used during

the backup, but place the restored files in the directory that is specified in the Alternate location text box. This option is useful when you do not want the current copies of the files overwritten.

- *Single folder*: Restores folders and files to the location specified in the Alternate location text box. Do not maintain the original directory structure, but rather place all of the files restored into the specified directory. This option is useful when you are looking for a few specific files that are buried or lost in a complicated directory structure.

7. Click Start Restore.

This opens the Confirm Restore dialog box which can be used to modify advanced options. I will not explore these options in this book.

8. Click OK to start the restore.

The Restore Progress dialog box will open, displaying the status of the file restore. When this process is complete there will be Close and Report buttons in the top-right of the dialog box.

9. Click Close.

The Restore process is now complete, and you can close the Backup Utility.

Restore points in Windows XP

With Windows XP, Microsoft introduced a new system for backing out of driver changes — *System Restore*. When critical operations, like updating drivers, are conducted on your system, Windows XP automatically creates a *system restore point*. This restore point can then be used to revert the system back to a previous state in the event of an immediate failure. You can also manually create your own restore points.

With a restore point saved, if your computer encounters problems related to changes in its system files or settings, you can use a restore point to quickly revert your computer to a previous state. It is like having an Undo button for your operating system.

Restore points are a replacement for the F8 boot option of Last Known Good Configuration, although Microsoft still supports the older and widely known Restore feature.

Manually creating a restore point or reverting your computer to a previous restore point is done via the System Restore Wizard, shown in Figure 3-3. The following sections show you how.



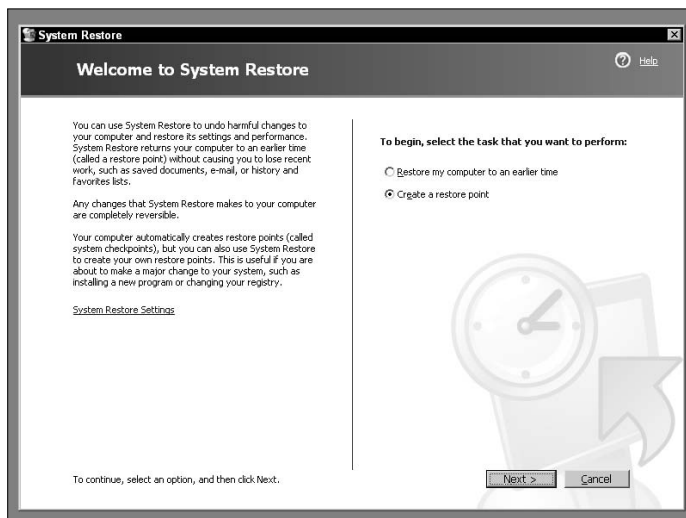


Figure 3-3:
The System
Restore
Wizard.

Manually creating a restore point

You can manually create a restore point from within the System Restore Wizard by following these steps:

1. Choose Start⇨All Programs⇨Accessories⇨System Tools⇨System Restore.

The System Restore dialog box opens and (refer to Figure 3-3) offers you two options:

- Restore My Computer to an Earlier Time
- Create a Restore Point

2. Select Create a Restore Point and click Next.

This moves on to ask you for a restore point description, which will make it easier to locate an appropriate restore point during the restore process.

3. Provide a description and click Create.

4. When the creation process is complete, click Close to exit System Restore.

In order for the system restore to work, Windows XP monitors all of the hard drives on your system. By default, up to 12% of your drive space holds system restore data. This system restore cache setting can be adjusted by choosing Start⇨Control Panel and double-clicking the System applet. Then click the System Restore tab. Figure 3-4 shows the dialog box and the settings for the C: drive which I accessed by selecting the drive and clicking the Settings button. You can adjust the system restore cache by moving the slider and clicking OK.

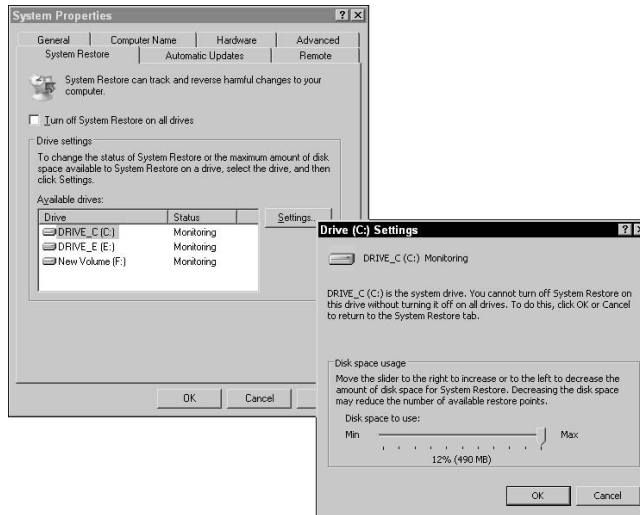


Figure 3-4:
Changing
the system
restore
cache size.

Reducing the size of the cache will reduce the disk space used by System Restore, but will also reduce the number of restore points that Windows XP will keep track of for you, which also reduces the age of your oldest restore.

If you only have one partition or drive on your computer, then you will be able to adjust the cache size without going into the settings for the disk.

Restoring your computer to a previous state

Minor and major disasters may require you to access a restore point to return your computer to working condition. The following instructions guide you through the process of restoring your computer to a previous state by accessing a system restore point:

1. Choose Start → All Programs → Accessories → System Tools → System Restore.

The System Restore Wizard (refer to Figure 3-3) gives you two options:

- Restore My Computer to an Earlier Time
- Create a Restore Point

2. Select Restore My Computer to an Earlier Time and click Next.

The wizard presents you with a list of restore points, as shown in Figure 3-5.

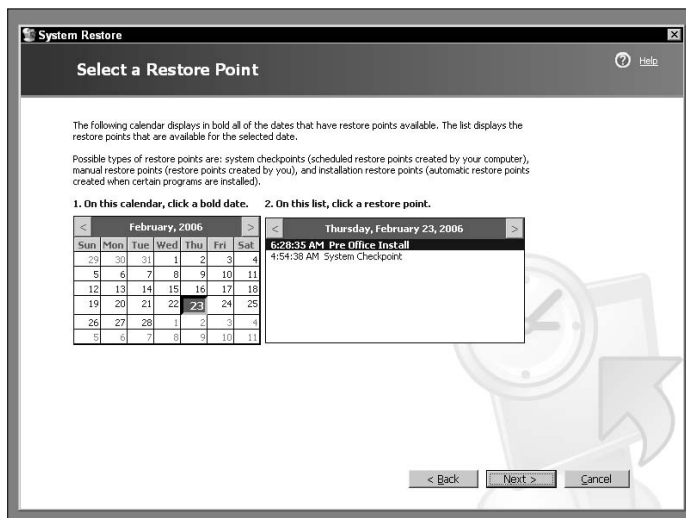


Figure 3-5:
Restoring a
system
restore
point.

3. Select the appropriate restore point and click Next.

You are presented with a dialog box warning you that you will lose recent changes on your system if you choose to restore.

4. Click Next.

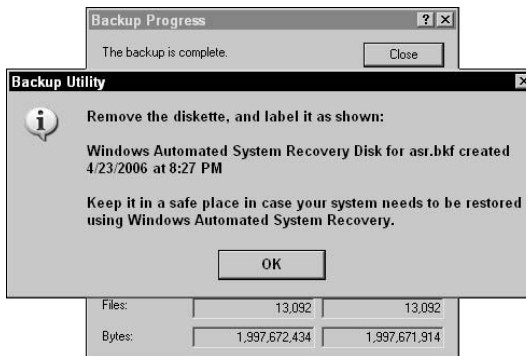
Your system reverts to how it was configured at the restore point you chose, and your computer reboots.

ASR in Windows XP

Before the development of the Automatic System Recovery (ASR) process, restoring a Windows computer by using the Backup Utility required a complete re-installation of Windows on the system, followed by a restore of the latest backup from backup media, which would be a file or tape. Now, the ASR process integrates a minimal Windows re-installation with the Backup Utility's restore process to reduce the total time it takes to get a system up and running.

In order to use the ASR features in Windows XP or Windows Server 2003, you need to first create an ASR backup. This backup is made by using the ASR Wizard which can be opened from within the Windows XP Backup Utility by choosing ASR Wizard from the Tools menu. In order to conduct the backup, you will require backup media and a blank floppy disk, which the backup process will ask for at the end of the ASR backup, as shown in Figure 3-6. The ASR backup does not perform a complete backup of your system; it only backs up the Boot partition and System State — any data on other partitions needs to be backed up and restored separately.

Figure 3-6:
In order to complete the ASR backup, you need a floppy disk.




If you lose your ASR floppy disk, you will not be able to restore that ASR backup.

After the backup is completed, you can restore the system by using the following process:

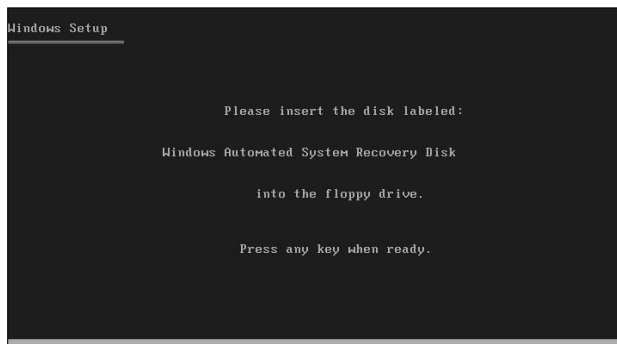
1. **Boot the system from the setup CD.**
2. **Press F2 when prompted to start the ASR process, as shown in Figure 3-7.**

Figure 3-7:
To start the restore process, you press F2 during setup.



3. **Insert your backup disk when prompted, as shown in Figure 3-8.**

Figure 3-8: You use the floppy disk you created during the backup to complete the restore.



The Automated System Recovery Wizard appears, as shown in Figure 3-9. The process proceeds automatically, or you can choose to go through it manually. It will perform a base OS installation onto the same partition that was backed up and then attempt to restore the backup that was conducted during the ASR process.

When the process is complete, your system should be in the same state that it was in during the ASR process.

Figure 3-9: The ASR restore process is automatic.



This process was designed to speed up the recovery time for systems that require a restore from backup. With the ASR complete, you can then restore data that is found on other partitions of your system. If the only problems that you encountered were related to your boot partition, then your job is complete.

Recovery and Rescue Methods

So far in this chapter, you have seen many different tools that you can use to repair your computer when it is in a non-booting state. This has included boot disks; Microsoft–provided tools like the Recovery Console; and the Windows Backup Utility. This section will look at manufacturer recovery tools, third-party tools, and special Windows installations.

Recovery CDs and Recovery Partitions

Many large computer companies, such as HP and IBM, provide a quick way to restore your computer to its default factory installation. In the past, this was done by using recovery CDs. When you unboxed your computer, you would get a set of bootable CDs that would contain some type of image of your computer. In the event that your computer became unusable due to a virus, hard drive failure, or other issue, you could quickly boot from the CDs and restore your computer to its factory state. One problem with this scenario is that it usually meant that you also lost all of your data if it was not otherwise backed up.

Over the years, with vendors prepackaging applications with their computers, the size of that factory image has gotten larger and larger, so that now it would take multiple CDs, or even multiple DVDs, to store all that data. This change in the preparation of the computers has caused some vendors to use a recovery partition instead. This hidden partition contains the tools needed to restore the image to the rest of the hard drive. If you use partitioning tools and delete the partition, then you lose your ability to restore the system to its factory default settings. The vendor will usually, however, give you the ability to back this partition up to CDs or DVDs. Usually, this partition is accessed by pressing one of the computer's function keys during boot-up.

Rescue CDs

A rescue CD is a bootable CD that includes tools that can be used to rescue files from a failing hard drive, or repair problems that have occurred with a Windows installation which made it non-bootable. Many third-party rescue CDs are available to the CompTIA A+ Certified Professional. Some rescue CDs are commercially available, but a wide variety of them are available for free. A quick Internet search for `Live CD` will reveal a wide range of these CDs. It seems that there are as many types of custom Live CDs as there are problems that you need to solve. With a little searching and experimenting, you can find one that meets your most common problems to add to your arsenal of tools.

Many rescue CDs run a version of Linux, which has been prepackaged onto a bootable CD or CD image which includes a variety of tools. Knoppix, running a version of Debian Linux, is just one of these Live Linux CDs that many people use as a rescue CD, partly because it has a wide selection of tools

and partly because it includes a graphical user interface in the form of the X Window System and KDE. With the variety of applications that are bundled on the Knoppix CD, and the fact that the system runs entirely off of the CD, Knoppix is a great way to try Linux without having to install it. The following offers a brief overview of some of the rescue options that are available to you when using Knoppix, and most of the other Live CDs offer a similar mix of features. You will have to try a few out to find one that works for you.



Conveniently enough, Wiley Publishing also publishes *Knoppix For Dummies*, by Paul G. Sery, which includes a complete Knoppix 4.0 distribution on DVD. So you can have your cake and learn how to eat it, too.

The Knoppix boot process starts off with a confirmation of your intention to boot into Knoppix, as shown in Figure 3-10. There are advanced boot options that you can access or get information about by pressing F2 or F3, but in most cases, the system will boot without additional information.

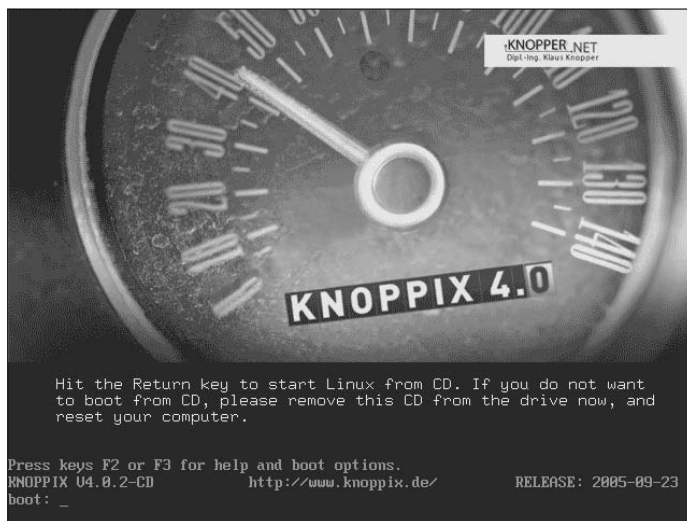


Figure 3-10: If your computer has unique hardware, you may need to use special boot options.

If you aren't familiar with Linux, here's a pointer: Linux identifies hard drives with letters and partitions with numbers, so in a typical IDE/ATA system, your four devices or drives are identified as `hda`, `hdb`, `hdc`, and `hdd`, with the `C:` drive usually being `hda1`. In the case of Knoppix, all of your local disk partitions are automatically mounted for read-only access and will appear on the Desktop. You can change the properties of the disk to allow writing to it; otherwise, you can browse and read files on your system. Figure 3-11 shows the properties of a Windows shortcut file that is found on `hdb1` in a `labfiles` directory.

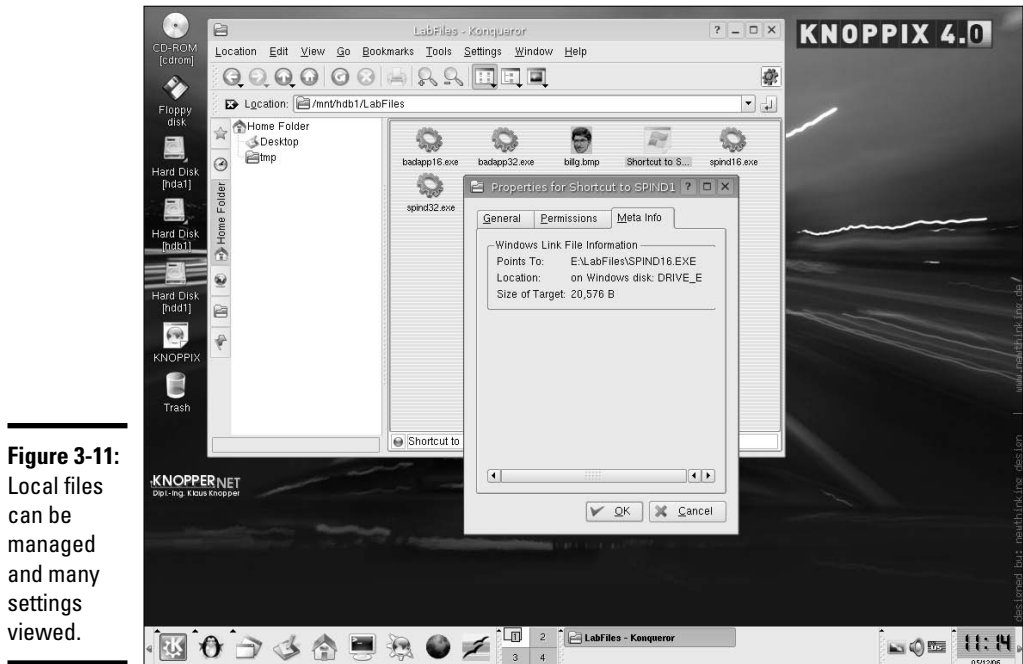


Figure 3-11: Local files can be managed and many settings viewed.

By using the tools on this Live CD, you will be able to verify many settings and even modify files that are found on your Windows installation. In Figure 3-12, the `boot.ini` file on `hda1` (the `C:` drive) is being modified with correct settings for the controller and partition.

Knoppix has several tools that you will likely find useful, such as `partimage` (partition image), which allows you to perform disk imaging tasks on your partitions, and `parted`, which is a command line disk partition editing tool. Figure 3-13 shows the interface of `QTParted` (a GUI front end for `parted`). This interface is reminiscent of Partition Magic. You can use `QTParted` to create, resize, delete, and move partitions around your drive.

If you need to copy files to or from a Windows server on your network, you can use tools like Samba Network Neighborhood to access shared folders on your servers. You can use this tool to copy network drives to a computer that doesn't have them or to rescue files from a hard drive that is very near failure and is no longer able to support booting. Figure 3-14 shows the shares that are available on the server `192.168.1.3`.

Figure 3-12: Editing local files to correct configuration problems.

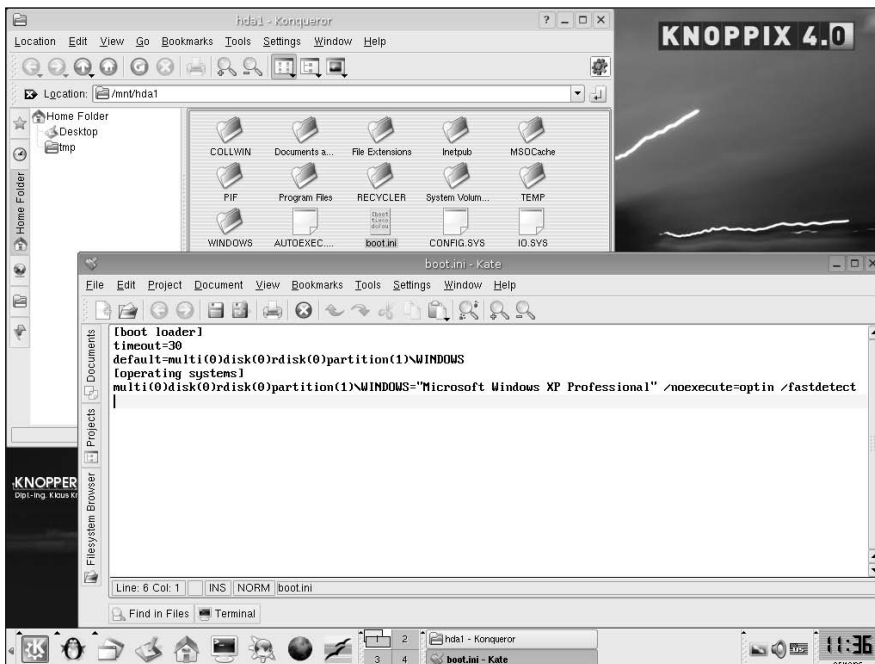
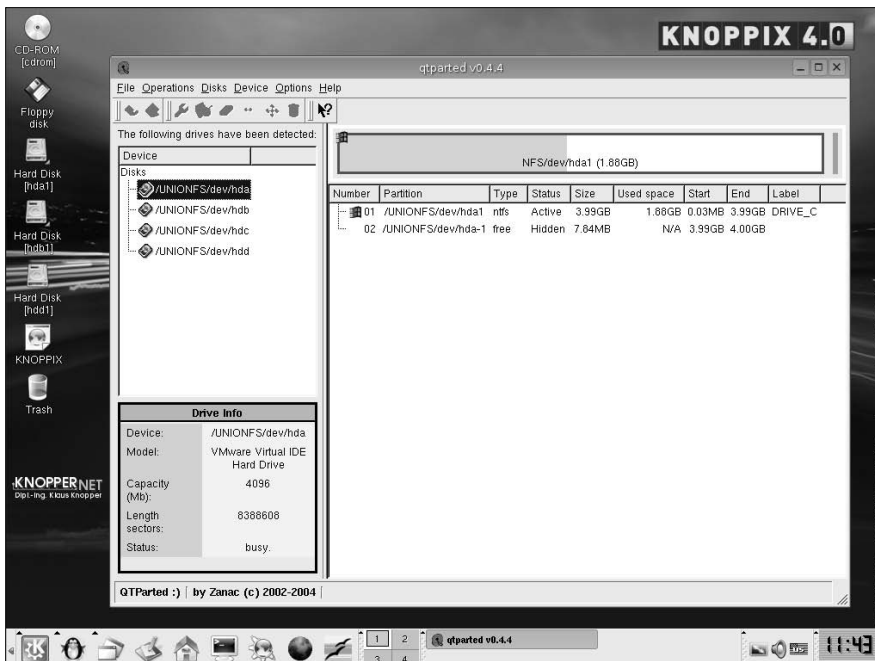


Figure 3-13: The disk partition management is easy to use, once you understand how Linux refers to disks and partitions.



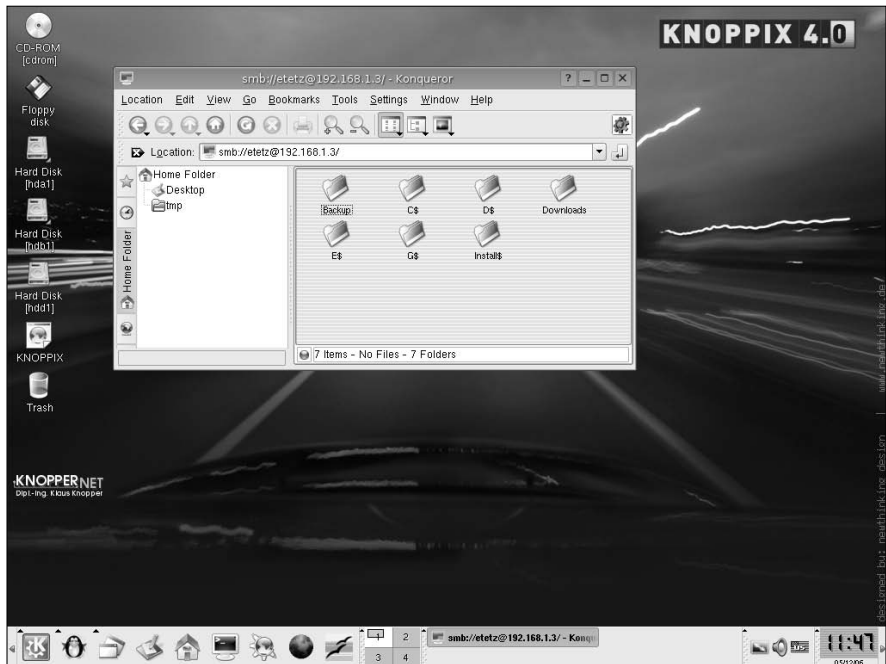


Figure 3-14: Files may be transferred to and from file servers.

Rescue Partitions or Rescue Installations

Many of the quickest ways to repair a Windows 2000 or Windows XP installation are actually built into Windows 2000 and Windows XP. To repair a Windows installation from a Windows installation, you need only a small amount of planning. If you leave some free, unpartitioned space on your hard drive, you can create a rescue partition. A *rescue partition*, or *rescue installation*, of Windows is simply another copy of Windows that is installed onto your computer. This second copy of Windows is not the default for booting, and is used only to repair problems to your normal installation of Windows. So if your default copy of Windows is installed on Drive C:, then your rescue install could be installed on Drive D: or Drive E:. The partition should be used only for the Windows installation, and not used to store data for your normal installation. After installing Windows in the second location, it will be automatically set as the default boot. To change the default booting installation of Windows, follow these steps:

1. Choose Start→Control Panel.

The Control Panel folder should open.

2. If your Control Panel folder is displayed in Category View, select Performance and Maintenance → System, or if your Control Panel folder is displayed in Classic View, double-click on System.

The System properties dialog box should now be open.

3. Click on the Advanced tab and then on Settings in the Startup and Recovery section.

The Startup and Recovery dialog should now be open.

4. Select the copy of Windows that you would like to be the default booting installation from the Default Operating System drop down menu, and click OK twice to close both of the open dialog boxes.

Getting an A+

This chapter helps you prepare for disasters by covering the following topics:

- ◆ Using the NTBackup.exe utility to protect your data.
- ◆ F8 boot options that can be used to correct startup problems.
- ◆ System restore points and how to use them to repair configuration problems.
- ◆ Recovery CDs and partitions and rescue CDs and partitions.

Prep Test

- 1** When enabling a logged boot of Windows 2000, what is the name of the log file that is created?
 - A ntbtlog.txt
 - B bootlog.txt
 - C btlog.txt
 - D ntlog.txt

- 2** What command would you use from the Recovery Console to attempt to repair the boot sector?
 - A listsvc
 - B disable
 - C systemroot
 - D fixboot

- 3** Which of the following files is not required on a Windows XP boot disk?
 - A ntldr
 - B ntbtlog.txt
 - C ntdetect.com
 - D boot.ini

- 4** A user bought a new video card with 128MB of RAM on it. He needs it to play the new version of his favorite video game. After installing the video card, he boots up his Windows XP computer and is overjoyed when he sees the Plug and Play Wizard detect the new hardware. When prompted, he provides the driver disk that came with his computer. After loading the driver for his new video card, his computer reboots. During the next boot, his computer halts and will not respond to any controls. What should he do next?
 - A Reboot his computer by using the power button. His computer will automatically boot to the Last Known Good Configuration.
 - B Reboot his computer by pressing Ctrl+Alt+Del and then pressing F8 when he sees the message to Press F8. He will then be able to choose Disable Advanced Video Options from the boot menu.
 - C Reboot his computer by using the power button. His computer will automatically display the boot menu and attempt to boot to Last Known Good Configuration without user intervention.
 - D Reboot his computer by pressing Ctrl+Alt+Del and then hold down F9 to enable Safe Mode.
 - E Reboot his computer by using the power button. When he sees the message to Press F8, press it and select either the Last Known Good Configuration or Enable VGA Mode.

- 5** Where does Windows XP store the boot log file, `ntbtlog.txt`?
- A The root directory of the first hard drive, typically `C:\`.
 - B The `C:\Windows` directory.
 - C The `C:\Windows\system32` directory.
 - D The `C:\recoveryerrors` directory.
- 6** What items are required to perform an ASR system restore? Choose all that apply.
- A The ASR floppy disk
 - B Your most recent backup
 - C Your ASR backup
 - D Your ERD floppy disk
 - E A Windows XP installation CD
- 7** When you boot the F8 option of Last Known Good Configuration, what action restores the configuration?
- A Deleting the recent drivers that have been installed. All driver locations are stored in `C:\Windows\inf`.
 - B Restoration of old registry keys from backup tape.
 - C Execution of `rollback.com`, to reset the correct configuration.
 - D Replacing the registry key `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet` with a backup copy.
- 8** What operating systems can create an ERD for the Emergency Repair Process? Choose two.
- A Windows 2000
 - B Windows 98
 - C Windows NT
 - D Windows XP
- 9** What tasks are completed by the Windows XP Emergency Repair Process? Select all that apply.
- A Inspect boot files
 - B Inspect driver configuration
 - C Inspect system files
 - D Inspect Registry
- 10** What type of backup will back up all selected files and clear the archive attribute of all files that are backed up?
- A Daily Copy
 - B Differential
 - C Copy
 - D Normal

Answers

- 1 A. The file is named `ntbtlog.txt`. See *"Booting into Safe Mode."*
- 2 D. `fixboot` is the command to repair the boot sector. Review *"Recovery Console."*
- 3 B. The only files that are required on a Windows XP boot disk are `ntldr`, `ntdetect.com`, and `boot.ini`. Check out *"Working with Boot Disks."*
- 4 E. Either of these two options, Last Known Good Configuration or Enable VGA Mode will allow him to boot his computer. If he chooses to boot the Last Known Good Configuration, then any changes that he made to the hardware configuration since the last boot will be lost. If he chooses the Enable VGA Mode, then he will be able to select a different driver to be used by the graphics controller or apply other software patches to get the system running. Safe Mode should be his next step on correcting the problem. Peruse *"Using F8 Options during Boot-up."*
- 5 B. The boot log file is stored in the Windows directory. For more information, see the section titled *"Tracking the boot with a bootlog."*
- 6 A, C, E. In order to complete an ASR restore, you need an ASR floppy disk, your ASR backup, and the Windows XP installation CD. After completing the ASR restore, you may want to restore a more recent backup to update any files on your boot drive that were changed since the ASR backup or to restore files from partitions other than the boot partition. The ERD is used when performing emergency repairs on Windows 2000 computers. Take a look at *"ASR in Windows XP"*
- 7 D. Peek at *"Using the Last Known Good Configuration."*
- 8 A, C. Only Windows 2000 and Windows NT can create an Emergency Repair Disk (ERD) to be used in the Emergency Repair Process. This disk is not required for Windows XP to complete the Emergency Repair Process. Examine *"Emergency Repair Disk (ERD)."*
- 9 A, B, D. The Emergency Repair Process does not inspect driver configurations. Check out *"Emergency Repair Process."*
- 10 D. The Normal backup will back up all selected files and clear the archive attribute on all files that are backed up. Review *"Understanding Backup and Restore."*

Book VIII

Networking

The 5th Wave

By Rich Tennant



Contents at a Glance

Chapter 1: Down to the Networking Basics	867
Chapter 2: Understanding Common Wireless Communications	907
Chapter 3: Networking the Operating System	927
Chapter 4: Configuring Internet Access.....	981

Chapter 1: Down to the Networking Basics

Exam Objectives

- ✓ Identifying the types of networks
- ✓ Understanding network topologies
- ✓ Working with network cables
- ✓ Becoming familiar with network architectures
- ✓ Accessing the network

The A+ Certification exams cover two areas of networking: networking theory/networking hardware and networking at the operating system level. This chapter focuses on the networking theory and the networking hardware area of the A+ exams. For the exam you are required to know popular terms and features of networking environments that you will encounter on the exams and in the real world — which is what this chapter will help you with. Networking at the operating system level is covered in Chapter 3 of this minibook.

Identifying the Types of Networks

A *network* is a group of systems that are connected together for the purpose of sharing data or sharing devices. This section provides an overview of the two major types of networks: *peer-to-peer* and *server-based* (client-server). I discuss the advantages and disadvantages of each type, as well as how to implement them.

Peer-to-peer networks

In a *peer-to-peer* (P2P) network, all systems connected to the network can act as clients or servers. A *client* is a system that makes a request for a resource or service on the network while a *server* is the system providing the resource or service. In this type of networking environment, all systems are considered equal with one another because they can all play the same roles on the network — either as client or server or as both client and server. The recommended number of systems in a peer-to-peer network

usually involves ten or fewer systems because of the lack of centralized administration. As a network administrator working in a peer-to-peer environment, you will constantly run from machine to machine to perform administrative tasks. Typically, a peer-to-peer network involves each system running a desktop operating system, such as Windows XP, to provide network functionality (see Figure 1-1).

In Figure 1-1, notice that system A provides a network resource — a printer — as does system D. This shows that system A is acting as both a server and a client, which is the purpose of a peer-to-peer network.

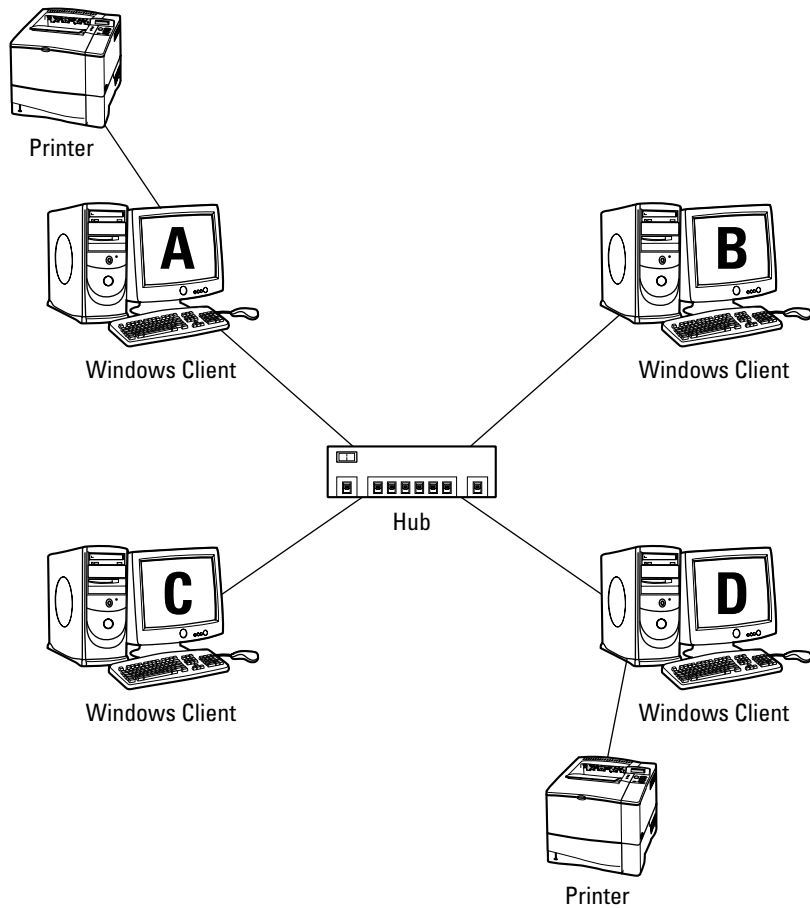


Figure 1-1:
Peer-to-peer
network
environ-
ments.

Because a central machine doesn't store files in a peer-to-peer network, your networking environment isn't based on the centralized administration approach. With *centralized administration*, you, as the network administrator, could perform network administration tasks for the entire network from one place. Looking back to Figure 1-1, you can see that because all four computers act as *peer servers* (meaning they are all acting as servers to one another), you need to do the administration on all four computers — a major disadvantage of peer-to-peer networking.

Some examples of the administration you must perform on each system are the creation of user accounts on each computer and the sharing of files and folders from each system. In Figure 1-1, for example, if you want Bob to log on to Computer A, you would create the `Bob` account on Computer A. At the same time, if you want Sue to log on to Computer B, you would create the `Sue` account on Computer B. Because the `Bob` account doesn't exist on Computer B, Bob can't log on to that computer, although he might be able to access the files on Computer B from Computer A. This leads to a *distributed administration* model because your work is spread out across multiple machines.

The major advantage of peer-to-peer networking is that you save money by not needing to purchase a central server, which can cost the company thousands of dollars in hardware and software. Keep in mind that with a peer-to-peer network, you save money not only in the hardware area, but you also don't have to purchase a separate network operating system. The *Network Operating System (NOS)*, required on a server-based network (discussed in the next section), is designed for networking services such as DHCP, Web, file, and print services, and allows the server to share its files and printers with clients on the network. The cost of the NOS, and the licenses to have clients connect to the server, is where a number of large companies spend most of their IT budget. Licensing is expensive!

Server-based (client-server) networks

Server-based networking, also known as *client-server* networking, is the networking model that most companies usually choose when there are ten or more workstations on the network. Unlike a peer-to-peer network, server-based networking uses a central machine (the server) that delivers network services to the workstations. Once again, these network services could be services such as file and print sharing, user account authentication, or Web services.

The benefit of a client-server configuration is that you can leverage centralized administration by performing the bulk of your work on the one server. For example, if you needed to create user accounts for each of the ten users, you would create the ten accounts on the one server, unlike in a peer-to-peer

network, where one account is created on each system. As the administrator of this network, you create all of the shared directories on the server along with user accounts so that the server may verify the credentials of a client who attempts to log onto the network. All users on the network connect to this server to save and retrieve files.

Tighter security is also a benefit of a server-based networking model. Creating a more secure environment is easier with a server-based network because your resources and user accounts are not spread across multiple machines. Looking at Figure 1-2, you can focus on the server because it contains the files, folders, and user accounts. When a user logs on to the network, the logon request is sent to the server, which verifies that the username and password are valid. After a user is logged on, the server allows the user access to resources, such as files and printers, that the user has permission to use. Figure 1-2 illustrates a server-based networking environment.

Notice in Figure 1-2 that the client systems connect to the server to access the printer. In this environment, all systems have a defined role — they are either a client or a server, but not both.

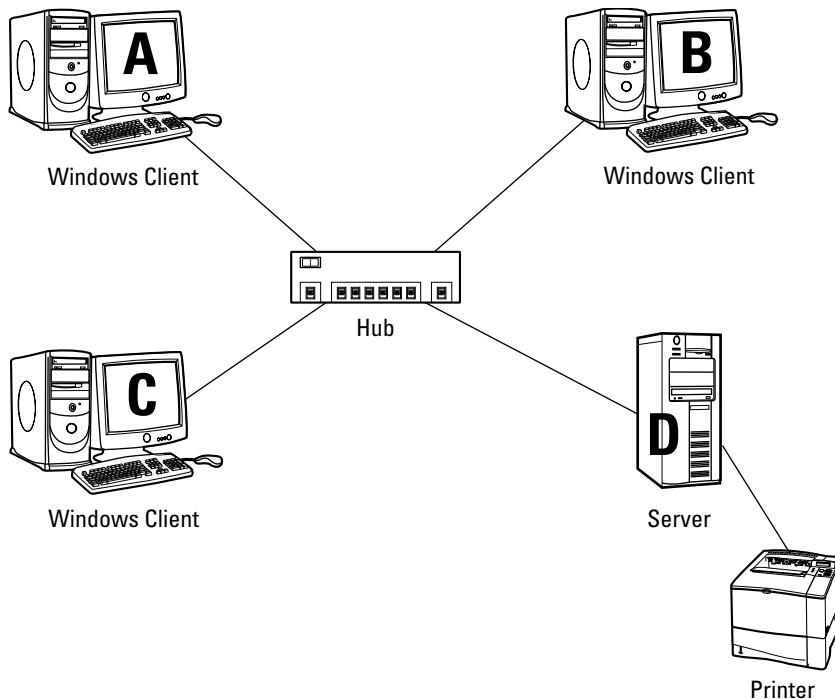


Figure 1-2:
Server-based
networking
environment.

The disadvantage of a server-based environment is the cost involved in purchasing the server hardware and the network operating system. When designing your networking model you want to ensure that you work with someone familiar with software licensing to ensure that you are getting the best bang for your buck!

Additional networking terminology

A set of terms you may hear when talking with other IT professionals about network concepts is *Local Area Network (LAN)* and *Wide Area Network (WAN)*.

A *LAN* is a network that typically involves one office building, or maybe even networked systems on one floor. The major point to remember when identifying a LAN is that there isn't a lot of distance between the systems on the network.

A *WAN* is a network environment that involves connecting two or more LANs. Each LAN typically covers its own building or office location. Companies normally like each office location to network with the other office locations. The connecting of all office locations creates the WAN.

Understanding Network Topologies

When building a network, it is important to understand some of the decisions that need to be made regarding the overall setup of the network. Building a network is like building a database: You have to understand the theory before you start the hands-on work.

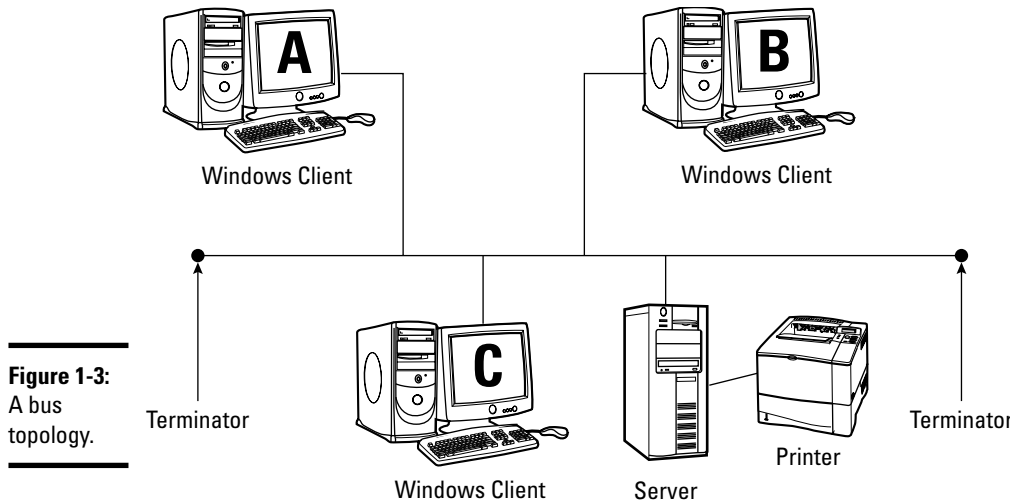
Topology refers in a general sense to *layout*, so a network topology defines the layout of the network. There are three basic network topologies: bus, star, and ring. The following sections discuss the different network topologies and their characteristics.

Bus

A *bus topology* uses a main wire (or trunk) to connect all network devices so that they can communicate with one another. The main trunk is fairly cheap to install but expensive to maintain. Figure 1-3 shows a diagram of a bus topology; notice that all systems are connected to this main cable length.

When a workstation sends data to another workstation in a bus topology, the data, in the form of an electrical signal, is delivered across the full length of the trunk. Each workstation looks at all data that runs along the trunk, and if the data is destined for a particular workstation, that workstation copies the

data to the memory on its network adapter. For example, Figure 1-3 can be used to demonstrate what happens when system A sends information to system B. The information runs along the trunk, and when it passes by system C, system C checks to see whether it is also a destination for the information; if not, system C ignores the data. The information continues down the wire and makes its way to system B. System B looks at the data to determine whether the data is destined for it; if so, system B copies the data and stores it in the memory on the network card. Note that because system B has made a *copy* of the data, the data is still on the wire. The data continues on the wire past the server and hits the terminator at the end of the trunk segment.



A *terminator* is a device that absorbs the electrical signal when it reaches the end of the network trunk. If there were no terminator at the end of the cable, the signal would bounce back in the other direction and collide with any new data being placed on the wire. So, to prevent this collision, the terminator grabs any signals that hit it and ensures it is absorbed off the wire.

In a bus topology, any break in the wire will create a non-terminated end and, thus, signal bounce results in the collapse of the entire network.



Star

One of the most popular types of network topologies today is the star topology. A *star topology*, shown in Figure 1-4, involves a central component, called a *hub* or *concentrator*, which connects all systems and is used to send the electrical signal to all connected systems.




Figure 1-4:
A star
topology.

With the star topology shown in Figure 1-4, if System A sends information to System D, the information first travels from System A to the hub. The hub sends the information through each port on the hub, as a result reaching each workstation connected to the hub. Each workstation is responsible for determining whether it is the data's intended destination. When System D receives the data, it checks the destination address of the packet, identifies itself as the recipient of the data, and then copies the data to the network adapter's memory. If the data is not destined for the system, the system simply discards the packet.

One of the major benefits of the star topology is that if a cable breaks, it doesn't take down the entire network like it would with a bus topology — only the workstation connected to the broken cable is affected. If the hub device breaks, however, the entire network fails. Note that the cost of implementing a star topology may be a little more than a bus topology due to the price of the hub device.

Ring

In a *ring topology*, each computer is connected to the next computer, creating a physical ring. Although ring topologies are not common today, you still see them in IBM's token ring architecture (see "Token Ring," later in this chapter). Figure 1-5 shows a ring topology.




Figure 1-5:
A ring topology.

In environments that use the ring topology, data is usually passed from workstation to workstation. Because data becomes distorted when it travels great distances, each workstation is responsible for reading the data, then regenerating the data and passing the information on to the next workstation. As with a bus topology, any break in the ring causes the entire network to fail.

Hybrid

A *hybrid topology* is a mixture of two or all of the three basic topologies. For example, you could use a bus topology as a main trunk, connect hubs to the main trunk, and then connect the systems to the hubs. Figure 1-6 shows an example of this type of hybrid topology.

To be more accurate, the configuration shown in Figure 1-6 could be called a *star-bus topology* — a star topology mixed with a bus topology. Today, the most popular topology in use is the star topology — or maybe even a hybrid topology using a star-bus layout.

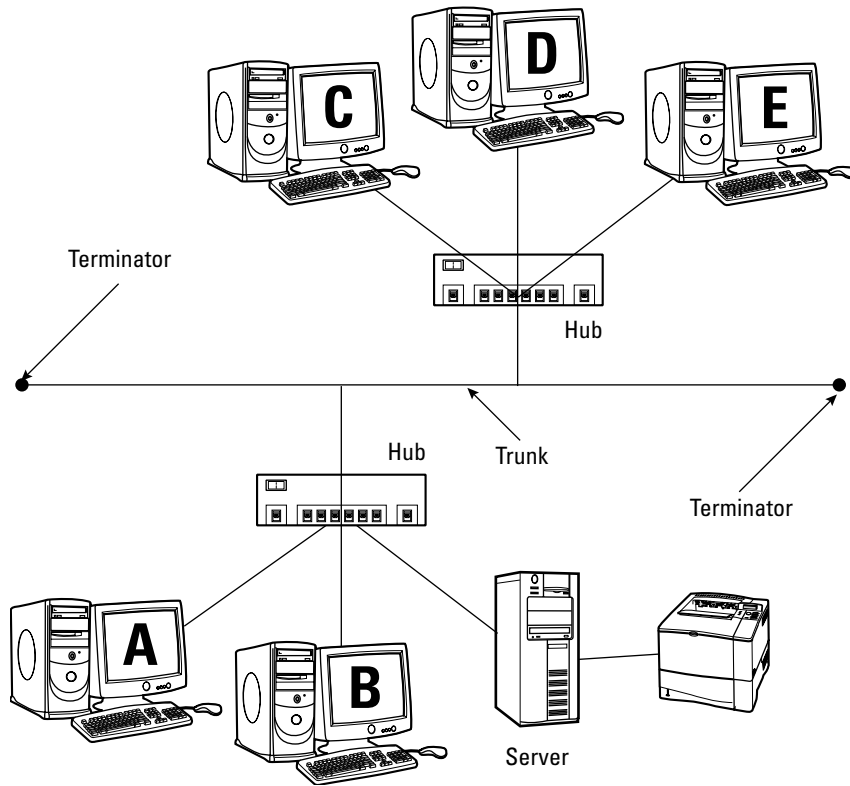


Figure 1-6:
A hybrid
topology.

Wireless

Today's networks allow more mobility out of the clients on the network by supporting wireless technologies. To implement a wireless solution, you build a wireless network that uses a wireless topology. A *wireless topology* typically involves a wired network with wireless clients connecting to the wired network through a *Wireless Access Point (WAP)*, a device that sends and receives signals to a wireless client in the form of radio waves (shown in Figure 1-7).

Notice in Figure 1-7 that the wireless client sends data to the wireless access point, which has a connection to the wired network. The WAP sends the wireless data to the destination system by sending the signal through the wired media.

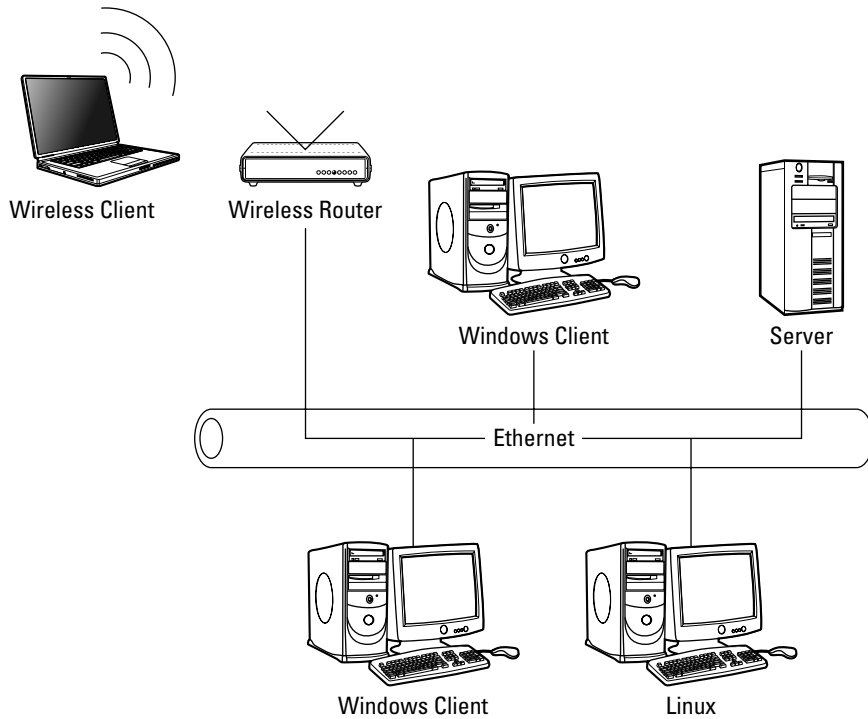




Figure 1-7:
A wireless
network
topology.

Connecting with Network Cabling

After you evaluate the different types of network layouts, it's time to connect all the network devices together, which means deciding the type of cabling you will use. The following sections discuss and evaluate the different types of cabling available for standard networks.

Twisted pair

Twisted pair cabling, which is inexpensive and easy to use, is one of the most popular types of cabling used today. It gets its name from the fact that it contains four pairs of wires twisted around each other inside the cable's outer jacket. The jacket is the term used for the outer covering of the cabling that is shown in Figure 1-8 and Figure 1-9. In addition, twisted pair cabling comes in two different flavors — *Unshielded Twisted Pair (UTP)* and *Shielded Twisted Pair (STP)* — shown in Figure 1-8 and Figure 1-9, respectively.



The two types of twisted pair cabling are fairly similar; the only difference is that STP cabling has an extra layer of insulation. The extra layer helps prevent interference from outside devices or cabling — interference that can distort the data traveling along the cable length.

Unshielded twisted pair cabling comes in a number of different flavors, called *grades* or *categories*. Table 1-1 lists the categories of UTP cabling, as well as their purpose and speed.

Table 1-1		UTP Category
Category	Purpose	Speed
Category 1 (CAT 1)	Voice only	
Category 2 (CAT 2)	Data	4 Mbps
Category 3 (CAT 3)	Data	10 Mbps
Category 4 (CAT 4)	Data	16 Mbps
Category 5 (CAT 5)	Data	100 Mbps
Category 5e (CAT 5e)	Data	1000 Mbps+
Category 6 (CAT 6)	Data	10 Gbps+

Because twisted pair cabling does not have the layers of shielding found in other forms of cabling, the data is pretty much unreadable — or the integrity of the data is questionable — after 100 meters. For this reason, twisted pair cabling has a maximum length of 100 meters.



For the exam, you are expected to know the speeds of the different categories of UTP cabling. Also, remember that twisted pair cabling has a maximum distance of 100 meters whether it is a shielded or unshielded twisted pair.

Twisted pair cabling uses a special type of connector to connect the cable to the system or network devices. This connector is similar to the one used to connect a telephone to a telephone jack. Network devices that use twisted pair cabling use the RJ-45 connector, while telephones use the RJ-11 connector. Figure 1-10 shows an example of an RJ-45 and RJ-11 connector.




Figure 1-10:
An RJ-45
connector
(left) and an
RJ-11
connector
(right).

Coaxial

Coaxial (coax) cabling is the type of cable you use for cable television. A copper wire in the center of the cable is responsible for transmitting information. Furthermore, the copper wire is protected by two levels of insulation and an exterior plastic covering, as shown in Figure 1-11.

Like UTP, coaxial cabling used for networking comes in different flavors — two to be exact. The first type of coax cable, called *Thinnet*, is only .25 inches thick, while the second type of coax, *Thicknet*, is .50 inches thick. Table 1-2 shows the difference between thinnet and thicknet.




Figure 1-11: Looking at coaxial cable and a coax cable with the BNC style connector.

Type	Coax Type	Maximum Cable Length	Diameter	Speed
Thinnet	RG-58	185 meters	.25 inches	10 Mbps
Thicknet	RG-8	500 meters	.50 inches	10 Mbps

Notice in Table 1-2 that the coaxial cable type is specified by what is known as a *Radio Grade (RG)*. There are a number of grade standards for coaxial cable, and each standard has a specific purpose and connector type that will work with that type of cabling. For example, RG-58-grade cabling, also known as thinnet, uses BNC connectors (discussed in more detail in the next section), while an RG-8 grade cable uses an AUI connector.

Another popular grade of coaxial cable is RG-6, which is the grade of coax used for television cable and which uses the F-Type connector to connect the cable to the device. Figure 1-12 displays coaxial cables along with an F connector and a BNC connector.

Connecting with thinnet

When using thinnet to connect to a workstation, you need to use a *British Naval Connector (BNC)*, which comes in a few forms. You will most likely encounter the plain old BNC and the BNC-T.




Figure 1-12:
An F-Type
connector
and BNC
connector.

The BNC connector connects thinnet cable to a networking device, such as a network card, using the barrel connector on the back of the network card, as shown in Figure 1-13.




Figure 1-13:
A BNC
connector
connecting
to a
workstation.

The BNC-T is shaped like the letter “T” and is used to continue the cable length and “T” off to connect a system to the network, as shown in Figure 1-14.




Figure 1-14:
A BNC-T
connector
used on
thinnet coax
cable.

Notice in Figure 1-14 that the BNC-T connects to a “metal barrel” type port on the back of the network card. If you don’t need to continue the cable length, and this is your last workstation, you are required to “terminate” the end with a terminator on the T connector, as shown in Figure 1-15.





Figure 1-15:
A BNC-T
connector
with a
terminator.

Connecting with thicknet

A system or device connected to a thicknet network uses an *Adapter Unit Interface (AUI)* port (shown in Figure 1-16), which connects the system to the thicknet cabling by using a transceiver known as a *vampire tap*.

Figure 1-16:
An AUI port.




The vampire tap gets its name by having small “teeth” that clamp to the cable and cut into the cable’s core, allowing the electrical signal to travel from the system to main network cable and beyond.

Fiber optic

Fiber optic cabling is one of the fastest types of network media available today. Fiber optic cabling is made up of a glass fiber core (optical fiber) surrounded by a layer of glass cladding insulation that is then covered with an outer covering, also known as a jacket. There are two fiber channels in fiber optic cable: one for sending information and the other for receiving information. Figure 1-17 illustrates a fiber optic cable.

Figure 1-17:
Fiber optic cabling.



Fiber optic cabling can carry a signal greater distances than twisted pair and coaxial cabling — fiber optic cabling can reach distances of 2 km or more. Fiber optic cabling can transmit information at speeds of 1 Gigabit per second (Gbps) and above. Because fiber carries data in pulses of light instead of electronic signals, it is impossible for the data to be corrupted by outside electronic interference.



You will be expected to know details regarding fiber optic cabling for the exam. Fiber carries data through pulses of light along its glass core and can reach distances of 2 km. Fiber optic cable transmits information at speeds that range from 100 Mbps to 10 Gbps.

There are two different implementations of fiber optic cabling — *Single-Mode Fiber (SMF)* and *Multimode Fiber (MMF)* — as explained in the following list:

- ◆ **Single-mode fiber (SMF):** SMF uses only one ray of light, known as a *mode*, to carry the transmission over great distances.
- ◆ **Multimode fiber (MMF):** MMF takes advantage of multiple rays of light, or modes, simultaneously. Each ray of light runs at a different reflection angle and is used to transmit data over short distances.



Fiber optic cabling uses a number of different types of connectors. The following list provides a few that you need to be familiar with for the A+ exam:

- ◆ **Straight Tip (ST):** The straight-tip connector is derived from the BNC-style connector but uses a fiber optic cable instead of the copper cabling that was used with BNC. Figure 1-18 shows an SC and ST connector.
- ◆ **Subscriber Connector (SC):** The subscriber connector is rectangular and is somewhat similar to an RJ-45 connector. Figure 1-18 shows an SC and ST connector.
- ◆ **Fiber Local Connector (LC) and Mechanical Transfer-Registered Jack (MT-RJ):** These are newer fiber optic connector types that resemble the registered jack and fiber SC shape. The MT-RJ is a small connector that is similar in appearance to an RJ-45 connector. The LC is similar in appearance to the fiber SC connectors and is the preferred connector for transmissions exceeding 1 Gbps because of its small form factor.

Figure 1-18:

Looking at the strands of fiber in fiber optic cabling along with the SC (bottom) and ST (top) connectors.



The primary disadvantage of fiber optic cabling is the cost of the implementation and the expertise that is required for the wiring.

Troubleshooting Networking Cables

Most network professionals use *cable testers* to test the cabling and verify that the cable is properly crimped and making contact with the networking devices.

For example, cable testers for CAT 5 cabling identify whether or not each of the eight wires has been crimped properly and identify problems such as wires being crimped in the wrong order.

Another problem you need to be aware of when running cable throughout a building is that twisted pair and coaxial cables are susceptible to outside interference from other electrical components. For example, you shouldn't run your networking cable alongside the electrical cabling because the electrical cabling could cause interference that could make the data on the network cable unreadable. For more information on cable testers and troubleshooting cabling problems, check out Chapter 2 in minibook 4.

Examining Network Access Methods

Network access refers to the different methods that computers use to place data on the network. This section discusses these methods and identifies the advantages and disadvantages of each.

CSMA/CD

One of the most popular types of access methods is *CSMA/CD*, which stands for *Carrier Sense Multiple Access/Collision Detection*. Understanding this term is easier if you break it down into its individual parts and examine each part in detail:

- ◆ **Carrier sense:** All computers on the network are watching, or sensing, the network for network traffic. If the network has data already on the wire, a system will wait till the wire is free of traffic.
- ◆ **Multiple access:** All computers on the network have equal access to the network at any given time. In other words, anyone can place data on the network whenever he or she chooses. Note, however, that workstations on the network will try not to place data on the wire at the same time the wire is transmitting other data because the two pieces of data will collide, destroying the data. That's why it's so important for workstations to "sense" the wire.

To summarize, “carrier sense multiple access” suggests that all workstations have access to the network and are watching the network to make sure it is clear of data before they send their information out.

- ◆ **Collision detection:** When two workstations send information out at the same time, the data will collide and be damaged in transit. When two workstations have data that has been involved in a collision, they resend the information out on the network at variable intervals to prevent the data from colliding again.

The nice thing about CSMA/CD is that the workstations decide when to send data, trying to prevent collisions. However, there is always the possibility that multiple workstations will send data out at the exact same moment the network is clear, resulting in data collision.

CSMA/CA

CSMA/CA, or *Carrier Sense Multiple Access/Collision Avoidance*, is similar to CSMA/CD except for one main difference: When a workstation senses that the wire is free, it sends out dummy data first instead of real data. If the dummy data collides with other information on the wire, then the workstation has avoided a collision with the real data; if the dummy data does not collide, then the workstation sends the real data. Like CSMA/CD, if the real data collides and does not reach its destination, it will be resent by the sender.

CSMA/CA is not a popular access method, but it has been used in AppleTalk networks.

Token passing

You may have attended meetings where a ball was passed around to indicate who had the authority to speak: In order to speak, you had to be holding the ball; if you didn't have the ball, you had to wait your turn.

Token passing is based on the same principle. A token, which is an empty piece of information running around the network from computer to computer, must be in a computer's possession before the computer can put data out on the wire.

When the token reaches a workstation, the workstation puts data out on the wire by filling the token with information and marking the token as being used. The token (with the information) is then released onto the network and travels toward its destination. Each workstation checks to see whether the token is destined for it when the token passes by. Each workstation regenerates the data and passes it to the next workstation.

When the token reaches the destination, the destination system reads the data and then sends a reply or confirmation message to the sending computer. When the sending computer receives the reply, it places a new, empty token on the wire.

Piecing Together the Network Architectures

A *network architecture* describes a network technology that uses a specific topology, cable type, and access method. This section describes the major types of architectures and their characteristics.



Many people use the term *architecture* to mean *topology*, stating that there are three main types of architectures: bus, star, and ring. However, bus, star, and ring are properly defined as topologies, not architectures. Be careful not to confuse the two terms. A *topology* defines the network layout, whereas an *architecture* is made up of a topology, cable type, and access method.

Ethernet

Ethernet is probably the most popular type of network architecture in use today. Ethernet is an example of a network architecture that comes in multiple flavors. If someone says to you, “I have an Ethernet network,” you are likely to ask, “What type of Ethernet?”

There are a number of types of Ethernet: 10BaseT, 10Base2, 10Base5, 100BaseT, and 1000BaseT being a few popular ones. All Ethernet architectures use CSMA/CD as the access method but each type of Ethernet uses a different topology or cable type. The following sections outline the differences among them.

10BaseT

10BaseT Ethernet is a network architecture that typically uses a star topology but may, in some cases, use a hybrid star-bus topology. 10BaseT networks primarily use CAT 3 UTP cabling, which transfers information at 10 Mbps. 10BaseT uses the CSMA/CD access method for putting information on the network.

When trying to remember the names of these network architectures, you can start by looking at the number at the beginning of the name — it indicates the transfer rate. In this case the transfer rate is 10 Mbps. You can then look at the tail end of the name to tell what type of cable the architecture uses. In this example there is a letter “T” at the end of the name, implying twisted pair. The word “base” in the middle means *baseband transmission*, which means that the signal takes up the entire width of the media when sending and receiving data.

10Base2

10Base2 Ethernet is a network architecture that uses a bus topology, but may sometimes use a hybrid star-bus topology. 10Base2 networks typically use thinnet coaxial cabling, which transfers information at 10 Mbps. 10Base2 uses CSMA/CD as its access method.

10Base5

10Base5 Ethernet is a network architecture that uses a bus topology. 10Base5 networks primarily use thicknet coaxial cabling, which transfers information at 10 Mbps. Like 10Base2, 10Base5 uses CSMA/CD as its access method.

10BaseFL

An old Ethernet architecture that ran at 10 Mbps and used fiber optic cabling was known as *10BaseFL*. You can remember that 10BaseFL used fiber cables by the *F* in the name. 10BaseFL uses CSMA/CD as the access method.

100BaseT

100BaseT Ethernet is one of the faster network architectures. It uses a star topology, but can also be found using a hybrid star-bus topology. 100BaseT networks primarily use CAT 5 UTP cabling, which transfers information at 100 Mbps. 100BaseT also uses CSMA/CD as its access method.

100BaseFX

100BaseFX is another Ethernet architecture that runs at 100 Mbps, but is different than 100BaseT in that it uses fiber optic cabling instead of UTP cable. 100BaseT and 100BaseFX are referred to as the “*Fast Ethernet*” standards. 100BaseFX uses CSMA/CD as the access method.

1000BaseT

There are a number of Gigabit Ethernet standards, which transfer information at 1000 Mbps. The first is *1000BaseT*, which uses UTP cabling. 1000BaseT uses CSMA/CD as the access method.

1000BaseSX

Another Gigabit Ethernet architecture is *1000BaseSX*, which uses multimode fiber optic cabling for short distances. You can remember this architecture by the *S* in the name, meaning *short* distances. 1000BaseSX uses CSMA/CD as the access method.

1000BaseLX

1000BaseLX is another Gigabit Ethernet architecture that runs at 1000 Mbps and uses fiber optic cabling. *1000BaseLX* is different than *1000BaseSX* by being the Gigabit standard using fiber that is designed for long distances. You can remember this standard by the *L* in the name, implying *long* distances. Since single-mode fiber optic cabling is designed for long distances, it is the type of fiber optics used for *1000BaseLX*. *1000BaseLX* and *1000BaseSX* use CSMA/CD as the access method.



Nowadays, you will probably purchase network cards that are marketed as 10/100 cards or maybe even 10/100/1000 cards. This means that the network card can be used with networking environments that transfer information at 10 Mbps, 100 Mbps, or 1000 Mbps.

Token Ring

Token Ring is a unique network architecture because it is completely different in design from Ethernet. As its name implies, *Token Ring* uses a ring topology with token passing as the access method. Although a Token Ring network can use almost any type of cabling, it typically uses CAT 5 UTP today.

Token Ring comes in two flavors: 4 Mbps and 16 Mbps. Token Ring also has a special name for the hub that connects all the workstations together — *MAU* (*Multistation Access Unit*). Because of this MAU, a Token Ring network appears to use a star topology, but internally, the MAU is connected as a ring, making a complete circle. Figure 1-19 shows a Token Ring MAU.




Figure 1-19:
A Token
Ring MAU.



For the exam, remember that all Ethernet environments use CSMA/CD as an access method while Token Ring architectures use token passing as their access method. Also remember that Token Ring was developed by IBM.

When IBM developed Token Ring architecture, the company also developed its own cable type and a proprietary connector known as the *IBM-type Data Connector (IDC)* or *Universal Data Connector (UDC)*. Figure 1-20 shows an IBM proprietary cable with an IDC connector.

Figure 1-20:
An IBM-
type Data
Connector
(IDC).



FDDI

The *Fiber Distributed Data Interface (FDDI)* network architecture uses token passing as an access method and fiber optic cabling as a cable type. FDDI also uses a ring topology like Token Ring does, but the difference is that with FDDI there are two rings, a *primary ring* and a *secondary ring*. The secondary ring is used for fault tolerance if the primary ring goes down, meaning that the second ring is only used if the primary ring fails.



Lab 1-1 will help you identify the different network architectures. Lab 1-1 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Understanding Network Protocols

To ensure that all of the networking components work with one another, there have been networking standards developed. If a company decides they want to develop something like a network card — they will ensure that the network card is following a standard so that it can communicate with all the other networking components. In this section, you find out about the different network standards that help each networking vendor develop networking components that function alongside other networking devices.

IEEE Standards

The *Institute of Electrical and Electronics Engineers (IEEE)* has developed a number of LAN standards that define the physical components of networking technologies. In these standards, the IEEE has defined such things as the way that network cards place data on the wire and the type of cabling used in different types of LANs. The LAN standards are defined by *Project 802*,

which was launched in February 1980. These 12 standards, shown in Table 1-3, define different networking architectures.

<i>Project</i>	<i>Description</i>
802.1	Internetworking
802.2	Logical Link Control (LLC)
802.3	Ethernet (CSMA/CD)
802.4	Token Bus LAN
802.5	Token Ring LAN
802.6	Metropolitan Area Network
802.7	Broadband Technical Advisory Group
802.8	Fiber Optic Technical Advisory Group
802.9	Integrated Voice/Data Network
802.10	Network Security
802.11	Wireless Networks
802.12	Demand Priority Access LAN

A few of the networking standards that you should be familiar with are 802.3, 802.5, and 802.11:

- ◆ **802.3:** This networking standard defines the Ethernet architecture, also known as CSMA/CD. This standard defines how data is placed on the wire. A system senses the wire to verify that it is free of data and then submits the data.
- ◆ **802.5:** This IEEE standard defines the Token Ring network architecture. This architecture uses the token passing access method, which is a set of rules that control how a system submits data on the wire.
- ◆ **802.11:** This IEEE standard defines wireless networking. Wireless networking has evolved over the last few years, and as a result, there are a few different types of wireless networks, wireless 802.11b and 802.11g being the two dominant standards:
 - *802.11b:* This wireless standard has a transfer rate of 11 Mbps.
 - *802.11g:* This wireless standard has a transfer rate of 54 Mbps and is backward compatible with 802.11b. This means that you can have an 802.11b *wireless access point (WAP)* and have a system with a more current 802.11g wireless network card, and both systems will still be able to communicate. They will use the lower transmission rate of the two, though!



For the exam, you don't need to memorize the entire list of standards. However, it is important to know that *Ethernet* is defined in project 802.3, *Token Ring* is defined in project 802.5, and that 802.11 defines the wireless standards because you will be tested on them.

For more information on wireless networking check out the next chapter!

Voice over IP (VoIP)

Voice over IP (VoIP) is a fairly hot topic these days. VoIP deals with allowing subscribers to the VoIP service to have telephone conversations over a TCP/IP network, such as the Internet. The benefit of such a service is that there are no long-distance charges — just your monthly subscription to the service.

VoIP not only allows for the transmission of voice but also other types of data, such as video. VoIP has become a popular method of communication that supports PC-to-PC communication, essentially using the PC as a phone.

Working with Network Devices

A computer network is a lot more than just a few computers and a tangle of cables — a computer network will need to have a number of different types of devices to connect the systems together or to connect the network to another network. This section identifies the different types of network devices popular in network environments today and covers the devices you will be tested on for the A+ exams.



Be sure to be familiar with these devices for the A+ exam and understand their benefits.

Network interface card (NIC)

The network interface card (commonly referred to as a *network card*) is responsible for connecting the computer or device to the network. More importantly, the network card on the sending computer is responsible for converting digital data into an electrical signal for copper or optical signal for fiber, that can be placed on the wire. The network card on the receiving computer is responsible for picking up the electrical signal and then converting it back to digital data that can be understood by the computer system.

Each network card has a unique address burned into the ROM chip on the card. This unique address is considered the *hardware address* of the network card because the manufacturer of the card burns it into the card. The address uniquely identifies your system on the network. The hardware address is also known as the *Media Access Control (MAC) address*. An example of a MAC address is 00-20-3F-6B-25-13.

Looking at the example of our MAC address, notice that the MAC address is made up of hexadecimal addresses — not your typical decimal numerals or binary values. Of the six groups of digits, the first three pairs identify the manufacturer of the network card, while the last three pairs are a unique set of digits assigned to a particular card built by the manufacturer. So, 00-20-3F, which appears at the beginning of the above mentioned MAC address, represents the manufacturer of the card, while 6B-25-13 identifies the card.

The MAC address of the sending and receiving system is stored in the header of the network packet that travels the network wire. When each system sees the packet traveling by it, it looks at the destination MAC address to decide whether the data is destined for it. If a system finds that it is the intended recipient for the data, it copies the data to its buffer — which is memory used to store the information while it waits to be processed. Figure 1-21 shows a network interface card.




Figure 1-21:
A network
interface
card.

Repeater

One major concern with cabling is the maximum usable length of the cable. For example, UTP cabling has a maximum length of 100 meters, while thinnet has a maximum length of 185 meters. The reason for putting a maximum distance on cable lengths is that the signal traveling along the cable becomes

too weak to read at the destination system by the time the maximum length is reached. The receiving computer is unable to read the information, so as a result does not acknowledge that it has received the data. When the sending computer doesn't receive an acknowledgement, it simply resends the data. This causes the information to be resubmitted and thus generates more network traffic.

One way to increase the distance of a cable length is to use a *repeater*. A repeater regenerates a signal so it can travel the extra distance. For example, the repeater shown in Figure 1-22 joins two lengths of thinnet coaxial cable. It joins the two cable lengths so that the signal can travel the distance from Computer A to Computer B. Note that this distance exceeds 185 meters, which is the maximum distance of thinnet. When the signal hits the repeater, the repeater rebuilds the signal so that it can travel another 185 meters.

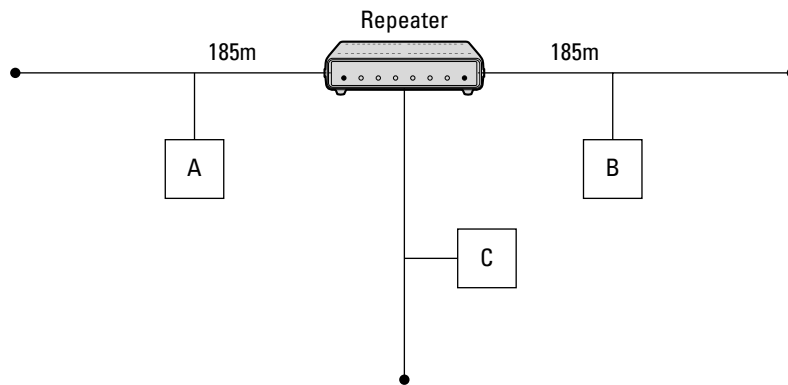


Figure 1-22:
A repeater
is used to
regenerate
the signal.

Bridge

Because all the data passing through a repeater is regenerated and sent to all parts of the network, a great deal of network traffic is generated that will affect the overall performance of the network.

To prevent this buildup of network traffic, you can use a bridge, shown in Figure 1-23. A *bridge* is a device that connects network segments together and also regenerates the signal (like a repeater). A bridge also filters the data so that it is sent only to the proper portion of the network, cutting down on network traffic and increasing overall performance.




Figure 1-23:
A bridge forwards network traffic to only the destination segment.

Figure 1-23 illustrates that when Computer A sends information to Computer C, the information travels along Segment 1 and eventually reaches the bridge. The bridge looks at its *bridging table* (a list of MAC addresses and corresponding network segments that runs in memory) to see which network segment Computer C exists on; in this example, it lives on Segment 3. At this point, the bridge forwards the information only to Segment 3, where Computer C resides, and not to any other segment, thus filtering traffic and cutting down on network noise.

Bridges increase performance on the network by filtering the network traffic, which as a result gives the network and all of its devices and applications more bandwidth to work with. The less network traffic, the less chance of collision and retransmission.

Router

A *router*, which is responsible for sending information from one network to another, is an important network device because, nowadays, most companies are connected to the Internet. When a computer on your network wants to send information to a computer on another network, your computer passes the information to your router. Figure 1-24 shows three different networks, each connected to the Internet by a separate router. All computers on Network A know that any information with an outside-network destination must be passed to the router because the router is the only device with a physical connection to the outside world.




Figure 1-24:
Routers
send data
from one
network to
another.

Gateway

A *gateway* is a unique network device responsible for converting information from one format to another. Think of a gateway as a translator between two different languages: as information passes from one side to another, the gateway “translates” the information to a format that can be understood on the other side.

As an example of a gateway, Microsoft has software called Gateway Services for Netware that can be loaded on a Windows Server. This software allows Microsoft Windows clients to connect to Novell networks without loading Novell client software by the Windows client sending the request to the Microsoft server (with the gateway software loaded). The Microsoft server will convert the data from Microsoft’s SMB (*server message block*) format to Novell’s NCP (*NetWare core protocol*) format and then send the request to the Novell server on behalf of the client — allowing Microsoft clients to communicate with a Novell environment without actually having a Novell client installed.

You may encounter questions on the exam in which you must identify the device based on a description. For the exam, remember that a *gateway* is a device or piece of software that translates data from one format to another.



Hub

A *hub* is a central device that acts as a connection point for all hosts on the network. A hub is a very basic device that passes all data that hits the hub to every port on the hub. This means that when a computer sends data to another computer, all systems will see the data on the network, although only the destination system for the data will process the data. Figure 1-25 shows a network hub.



Figure 1-25:
A network hub with eight RJ-45 ports and one BNC port.

As the number of hosts on the network grows, you can “cascade” or connect one hub to another. Any data that reaches a port on a hub will be sent to all ports on all connected hubs, which could congest the network with traffic.

Switch

A network *switch* is a device that looks similar to a network hub but differs in the fact that the switch does not forward the data to all ports like a hub would. Instead, the switch sends the data only to the port that the destination system resides on. Figure 1-26 shows a network switch.

Switches can dramatically increase network performance because they filter the traffic by sending the data only to the destination port on the switch instead of to all ports on all hubs.

Figure 1-26:
A 24-port
network
switch.



Wireless access point (WAP)

A very popular network component today is a *wireless access point (WAP)*. A wireless access point is typically connected to a wired network and is responsible for accepting data from wireless clients and then passing that data to systems on the wired network. The wireless access point can also receive information from the wired systems and then send that information to the wireless systems.

You can find a number of popular brands of wireless access points, such as Linksys, D-Link, and NetGear. Wireless access points that include additional features, such as firewall capabilities, are known as *wireless routers*. Figure 1-27 shows a wireless router used by home and small office networks.

Figure 1-27:
A four-port
wireless
router also
contains a
WAN link
for the
Internet
connection.



The wireless router has an antenna that collects the radio waves that carry the data from the wireless client. The wireless router also has a WAN (wide area network) port on it so that you can connect your Internet cable into it and share the Internet connection with all systems on the network. The WAP in Figure 1-27 also has four additional RJ-45 ports to connect four wired systems.

Understanding Communication Methods

Different network devices, such as network cards, support different methods of communication. The three major communication methods in the computer world are simplex, half-duplex, and full-duplex:

- ◆ **Simplex:** A device that supports *simplex* communication can deliver information in only one direction. Typically, the device can send information but it is unable to receive information.
- ◆ **Half-duplex:** A device that supports *half-duplex* communication can deliver information in two directions, but not at the same time. A typical example of a half-duplex device is a walkie-talkie or CB radio, which enables you to send and receive information but not at the same time. When you use a CB radio, you say “over” to let the person on the other end know that you are done talking and that it is the other person’s turn to send information.
- ◆ **Full-duplex:** A device that supports *full-duplex* communication is one that can send and receive information at the same time. Full-duplex is similar to talking on the phone — you are free to speak and to listen at the same time.

When purchasing networking or other types of devices for your computer, it is important to know whether you are buying a simplex, half-duplex, or full-duplex device to avoid problems later. For example, a musician friend of mine wanted to record his own material on the computer by using recording software. After recording the first track, he had a problem recording a second track because his sound card was only a half-duplex device. He wanted to listen to the first track while playing along and recording the second track, which was impossible with his half-duplex sound card. He needed to have the sound card send and receive information at the same time, something that a half-duplex device can’t do. The solution: get a new, full-duplex sound card.

Ways to Network a Computer

If you want to build a small network with four computers in the home and have each system connect to the Internet, the best thing to do is to buy one of those home routers provided by D-Link or Linksys. Once you buy the home

router, which is also a 4-port switch, connect the four computers to the ports on the router and then plug the Internet connection into the WAN port. All systems are now connected together and also to the Internet. But what if you need to connect two systems together and it is not possible to connect them by using a hub or switch? You need to be aware of the different and varied methods used to network two computers, including the following:

- ◆ **Installing network cards**
- ◆ **Using serial or parallel ports**
- ◆ **Using infrared ports**

The following sections describe the different computer networking methods.

Network card

The first task when using network cards to network two computers is to find out what type of network card you need for each system. You will have to find out whether the network card will be an ISA, EISA, or PCI device; you can do this by opening up the system and looking at the different expansion slots that exist, or you can look at the documentation for the system.

After you have purchased the network card and installed it into your system, you have to load the driver for the network card. If you're lucky and the network card is a Plug and Play device, then the driver might load automatically for you, or you might be prompted by the operating system to provide a manufacturer diskette that contains the driver for the device.

After you install the driver, you may have to spend some time troubleshooting the device because of resource conflicts with other existing devices. Once again, if you have a Plug and Play device, this step will probably not happen because the resources will be assigned dynamically by the operating system.

Serial and parallel ports

You may also network two computers by using the serial or parallel ports of both systems. Using standard parallel or serial ports for the purpose of networking two devices can be a lot slower than networking two computers via a network card. You will connect the parallel ports with a *laplink cable*, and you will use a *null modem cable* for serial ports (RS-232 ports).

There are actually two networking methods that use serial ports:

- ◆ You can connect a modem to each computer's serial port and have the modems use the phone lines as the cabling.
- ◆ You can connect two computers directly together by connecting the serial ports with a null modem cable.



If you try to connect two computers together with a normal serial cable, the *TD* (*transmit data*) wire on one computer will be connected to the TD wire on the other computer, and the *RD* (*receive data*) wire on one computer will be connected to the RD wire on the other computer. No communication will occur in this situation. What you need is the TD wire on one computer connected to the RD wire on the other computer. A null modem cable is designed to do this, crossing the sending and receiving wires.

Infrared port

Many systems today, especially laptop computers, have built-in *infrared ports* that can connect to other devices on the network. For example, you could print to a printer without using a parallel cable by using infrared technology.

Infrared technology uses an infrared light beam to carry data between devices. It typically requires clear line of sight — there must be a clear path between the two devices. Infrared is limited in distance to about 100 feet and can transfer information up to 10 Mbps.



Lab 1-2 will give you some practice creating a small office or home network. Lab 1-2 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Getting an A+

This chapter introduces you to a number of key concepts in network terminology. Understanding networking and how to troubleshoot the network has become an important skill for an IT professional, and you will be asked a few networking questions on the A+ exam. Remember the following points when preparing for the exam:

- ◆ There are three types of cables used on networks: *twisted pair*, *coax*, and *fiber optic*.
- ◆ There are three basic types of network layouts, known as topologies: *bus*, *star*, and *ring*.
- ◆ An access method determines how a system places data on the wire. Two popular access methods are *CSMA/CD* and *token passing*.
- ◆ *Ethernet* is the most popular network architecture used today. Other examples of network architectures are *Token Ring* and *FDDI*.
- ◆ There are a number of popular networking devices — some of the most popular devices are hubs, switches, and routers.
- ◆ *Switches* filter network traffic by sending the data only to the port used by the destination system. *Routers* are responsible for sending data from one network to another.

Prep Test

- 1** What category of UTP cabling transmits data at 16 Mbps?
 - A Category 2
 - B Category 3
 - C Category 4
 - D Category 5

- 2** Which of the following defines simplex communication?
 - A Allows information to be sent and received, but not at the same time
 - B Allows information only to be sent
 - C Allows information to be sent and received at the same time
 - D Allows information to be sent, but only after dependent information is received

- 3** What is the transfer rate of Category 3 cabling?
 - A 2 Mbps
 - B 10 Mbps
 - C 16 Mbps
 - D 100 Mbps

- 4** What is the maximum distance of a thinnet segment?
 - A 100 meters
 - B 500 meters
 - C 250 meters
 - D 185 meters

- 5** What is the recommended number of users for a peer-to-peer network?
 - A Fewer than 100
 - B More than 100
 - C Fewer than 10
 - D More than 10

- 6** What access method is used for Ethernet?
 - A Token passing
 - B CSMA/CA
 - C Twisted pair
 - D CSMA/CD

- 7 Which of the following defines half-duplex communication?**
- A Allows information to be both sent and received, but not at the same time
 - B Allows information only to be sent
 - C Allows information to be sent and received at the same time
 - D Allows information to be sent, but only after dependent information is received
- 8 What category of UTP allows for data transfer at 4 Mbps?**
- A Category 2
 - B Category 3
 - C Category 4
 - D Category 5
- 9 Which topology must be terminated at both ends to prevent signal bounce?**
- A Ring
 - B Star
 - C Bus
 - D Coaxial
- 10 What access method best describes CSMA/CA?**
- A Data is placed out on the wire; the sending workstation detects whether there is an error and retransmits if there is.
 - B A token runs around the network; when a computer wishes to send data out on the network, it fills the token with information.
 - C Dummy data is placed on the wire; if the dummy data collides with other information, then the real information is not transmitted. If the dummy data does not collide, then the real data is delivered.
 - D Dummy data is placed on the wire; if the dummy data collides with other information, then the real information is transmitted, but if the dummy data does not collide, then the real data is withheld.
- 11 Which of the following best describes a client-server environment?**
- A All users on the network connect to one another for the purpose of file sharing.
 - B All users on the network connect to a central server and access resources on that central server.
 - C All users on the network connect to one another for the purpose of printer sharing.
 - D Each user accesses only one other user's computer.

12 What is the maximum distance of UTP cable?

- A 100 meters
- B 185 meters
- C 250 meters
- D 500 meters

13 Which of the following best describes token passing?

- A Data is placed out on the wire; the sending workstation detects whether there is an error and retransmits if there is.
- B A token runs around the network; when a computer wishes to send data out on the network, it fills the token with information.
- C Dummy data is placed on the wire; if the data collides with other information, then the real information is not transmitted. If the dummy data does not collide, then the real data is delivered.
- D Dummy data is placed on the wire; if the data collides with other information, then the real information is transmitted. If the dummy data does not collide, then the real data is withheld.

14 Which of the following sends data in the format of pulses of light?

- A Fiber optic
- B 100BaseT
- C 10BaseT
- D Coaxial

15 What is the maximum length of 10Base5 cabling?

- A 100 meters
- B 185 meters
- C 250 meters
- D 500 meters

16 Which of the following best describes full duplex communication?

- A Allows information to be both sent and received, but not at the same time
- B Allows information only to be sent
- C Allows information to be sent and received at the same time
- D Allows information to be sent, but only after dependent information is received

17 What Ethernet architecture transfers information at 100 Mbps and uses Category 5 cabling?

- A 10BaseT
- B 10Base2
- C 10Base5
- D 100BaseT

18 What type of connector connects a Category 5 cable to a network card?

- A BNC Barrel
- B RJ-45
- C RJ-11
- D BNC

19 What is the maximum distance of fiber optic cabling?

- A 100 meters
- B 185 meters
- C 500 meters
- D 2 kilometers

20 What type of connector connects a 10Base2 cable to a network card?

- A AUI
- B RJ-45
- C RJ-11
- D BNC

Answers

- 1 C. Category 4 UTP cabling transfers data at 16 Mbps. *See "Twisted pair."*
- 2 B. Simplex devices deliver information in only one direction. *Review "Understanding Communication Methods."*
- 3 B. Category 3 cabling transfers information at 10 Mbps. *Check out "Twisted pair."*
- 4 D. The maximum distance of a thinnet segment is 185 meters. *Peruse "Coaxial."*
- 5 C. The recommended number of computers in a peer-to-peer network is 10 or fewer. *Take a look at "Peer-to-peer networks."*
- 6 D. CSMA/CD is the access method that is used in all Ethernet environments. Token passing is used in Token Ring architectures, and CSMA/CA has been used in AppleTalk networks. *Peek at "Ethernet."*
- 7 A. Half-duplex devices allow you to send and receive information, but not at the same time. *Look over "Understanding Communication Methods."*
- 8 A. Category 2 UTP cabling transfers data at 4 Mbps. *Study "Twisted pair."*
- 9 C. A bus topology must have all loose ends terminated to prevent signal bounce. *Refer to "Bus."*
- 10 C. With CSMA/CA, your computer avoids a collision with the real data by placing dummy data on the wire first. If the dummy data collides, your computer knows that it isn't safe to send the real information; if it doesn't collide, then the real data is sent. *Examine "CSMA/CA."*
- 11 B. Client-server environments are implemented for the purpose of centralized administration and security. It is much easier for an administrator to control resources if he is sitting at one computer and all users connect to that one computer. *See "Server-based (client-server) networks."*
- 12 A. The maximum distance of UTP cable is 100 meters. *Review "Twisted pair."*
- 13 B. With token passing, an empty token runs around on the network, and when your computer wants to submit information on the wire, it fills the token with the information and releases the token onto the network. *Check out "Token passing."*
- 14 A. Fiber optic cables send information in pulses of light on cabling with a glass core. *Peruse "Fiber optic."*

- 15 D. 10Base5 is thicknet that has a maximum distance of 500 meters. UTP cabling has a maximum distance of 100 meters, while thinnet has a maximum cable distance of 185 meters. *Take a look at "10Base5."*
- 16 C. Full-duplex devices can send and receive information at the same time. *Peek at "Understanding Communication Methods."*
- 17 D. 100BaseT Ethernet architectures can transfer information at 100 Mbps. *Look over "100BaseT."*
- 18 B. RJ-45 connectors connect UTP cabling to the network card. *Study "Twisted pair."*
- 19 D. Fiber optic cabling has a maximum distance of about 2 km. *Refer to "Fiber optic."*
- 20 D. A BNC connector connects a 10Base2 cable to a network card. *Examine "Connecting with thinnet."*

Chapter 2: Understanding Common Wireless Communications

Exam Objectives

- ✓ Identifying characteristics of Bluetooth, 802.11, infrared, and cellular networks
- ✓ Optimizing features of 802.11 wireless networks
- ✓ Identifying security settings for 802.11 access points and wireless networks
- ✓ Connecting to 802.11 wireless networks

It was once the domain of cutting-edge techno-geeks, but today, wireless technology is for everyone. Wireless technology is now entering all aspects of our lives, so as a CompTIA A+ Certified Professional, you will need to deal with wireless connectivity. This chapter gives you an overview of the technologies that you may need to deal with.

The main types of wireless devices that you will see in this chapter are infrared, Bluetooth, cellular networking, and 802.11 wireless networking.

Understanding Infrared Devices

Out of the wireless technologies that are covered in this chapter, the technology that has been around the longest and is likely the most stable in its development is infrared. *Infrared* uses light beams in the infrared spectrum, which is beyond the visible light spectrum. You likely use infrared technology daily if you have a remote-controlled television, VCR, or DVD player; infrared technology drives almost all of these small remote controller units. The technology, when used with computerized equipment, follows the standards of *IrDA* (*Infrared Data Association*). More information about the association and the Infrared standards can be found on their Web site at www.irda.org.

Infrared is a line-of-sight technology, which means you need a direct, unobstructed view between your transmitter and receiver. You may have noticed this requirement when using your television remote, and this requirement applies to all computer equipment using infrared. Because it is line of sight, it has obvious limitations as to where it can be used. Any objects obstructing the line of sight will prevent data transmissions. The benefit of line of

sight is that you will not have interference from or interference with areas that are outside of the line of sight, such as devices in the next room.

The goal of infrared networking was for short range (less than a meter), direct line of sight (+/-15 degree cone), and low speed (technologies between 16 Mbps and 2.4 Kbps). There are many different specifications, each with a different speed specification. One of those specifications is used for open office networking, a technology which has been replaced by 802.11b and involved infrared repeaters placed around an office. Currently, you will see infrared used for printers and personal digital assistants (PDAs).



Most tasks that were handled by infrared technology have switched over to using newer 802.11 networks and Bluetooth technologies.

Working with Wireless Networks

“Why do I need a wireless network?” has been a question that people used to regularly ask. This question does not appear on as many people’s lips these days. The biggest reasons people want wireless networks these days include mobility and not needing to deal with cabling issues.

When wireless networks started, they had access speeds of 1–2 Mbps; current standards have network speeds of 100 Mbps and beyond. At one time, wireless networking equipment was a premium addition to a network, but this equipment has become a commodity and is inexpensive enough that anybody can start up a wireless network.

Wireless networking should not be used for a primary network because of its limitations (see “Troubleshooting Issues,” later in this chapter). If the user knows what the limitations are, the network should be considered a best-effort network, and as such, it should not be expected to be available or to function at full speed at all times. As a secondary network, wireless networking offers a great deal of flexibility for getting to your data. Any devices that exist on your normal wired network can be present on your wireless network, including firewalls, servers, and printers.

Components of a wireless network

The two main components of wireless networks are clients and access points. Clients are computers that have *wireless network cards*; these cards serve the same purpose as wired network cards, but without the wires. *Access points* act as consolidation point for multiple wireless clients, and have a connection to a traditional wired network. Some new access points will allow you to link multiple access points together through wireless links, extending the range and coverage of wireless network, without having all access points connected to a wired network, thereby saving on wiring costs.

Many laptops now include wireless network cards as an integrated feature, and wireless network cards are available for desktop computers, connecting through PCI, PCI Express, or USB. It is great that many computers have wireless cards integrated in them, but be aware — as new technology arrives on the scene, you will likely want to use an expansion option (PCMCIA slot or USB port) to add an updated wireless card to your computer.

In the past, wireless networks commanded a premium price, but in the current market, wireless networking components have become a commodity. You now have several choices which are below \$100. The difference in price between the inexpensive units and expensive units is based on the strength of the radios, the features that are available on the unit, and its brand name.

If you are working with a wireless network with an access point, then the network client runs in Infrastructure mode. *Infrastructure mode* has access points making up the underlying network infrastructure. If the network is composed only of clients joined together to form a network, then the network is in Ad-hoc mode. *Ad-hoc mode* is usually used only when people need to exchange files or work together in a location that does not have an access point. This type of situation often happens in boardrooms when people are working on a project for a couple of days. In the past, these boardroom users would ask the network administrator for a 4–8-port mini-switch or hub which they would have used to share files; but they now have the ability to set up an ad-hoc wireless network. Both of these wireless network types are shown in Figure 2-1.




Figure 2-1: Wireless networks are usually Infrastructure mode, but may also be set up in Ad-hoc mode.



Most wireless networks will be composed of access points and wireless clients, so unless the question specifies *peer-to-peer* or *ad-hoc*, access points should be expected to be on the network.

Connecting to a network

Every wireless network is identified by an *SSID* (Service Set Identifier, which is used by all of the clients on the network, identifying them as members of the network. The SSID is a 32-character, case-sensitive name. Every manufacturer sets a default SSID, which for the purpose of security and to reduce confusion with neighboring access points, should be changed to a unique name. In order to connect to a wireless network, you need to know the SSID.



Some manufacturers refer to the SSID as a BSSID (Basic Service Set Identifier) or an ESSID (Extended Service Set Identifier). The 802.11-1999 wireless network standard defines an ESSID as a set of access points using the same SSID and channel, and operating as a single BSSID to their wireless network clients.

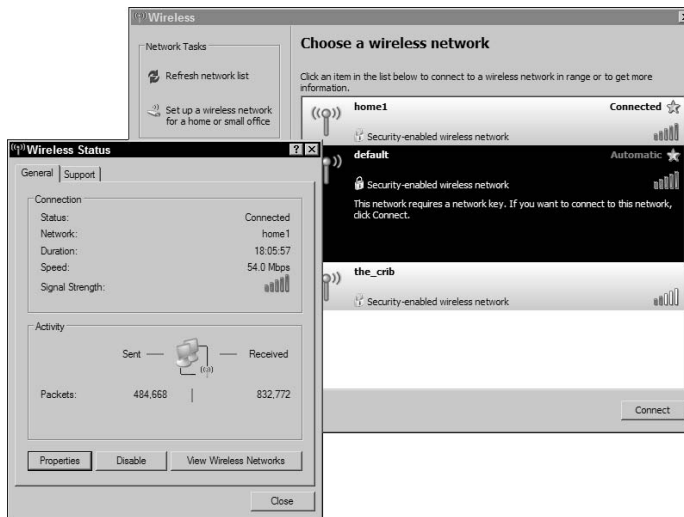
After you have your wireless network card installed in your computer, you have to decide whether you are going to use Windows XP's Windows Zero Configuration or the custom software that comes with some wireless network cards. If you are using the default Windows software, you can open your Network Connections folder and double-click your wireless network card to open the Wireless Status dialog box. This gives you information about your current wireless connection, including connection speed, SSID, and signal strength.

If you are not connected to the network that you want to be connected to, then you can click the View Wireless Network button, which shows you all networks in your area that are broadcasting their SSIDs. To configure a connection to one of the networks, simply select the network and choose the Connect button. If you need to enter a security key, then you will be prompted to enter the required information. Both the Wireless dialog box and the Wireless Status dialog box can be seen in Figure 2-2.

Wireless standards

There are several standards for wireless networks, and the performance and security features for wireless networks are constantly improving. What's more, because of constant improvements, new standards continually emerge. Gordon Moore of Intel predicted that the components in a processor's integrated circuit would double every 24 months. This statement has since been named Moore's Law, and has been applied to many areas of the computing industry. As Moore's Law is applied to more sections of the computer industry, wireless is just another section of the industry that falls to those general rules, with speed doubling every two years.

Figure 2-2: Wireless configuration is easy and straightforward with Windows XP.



The main wireless standards currently in use are 802.11a, 802.11b, and 802.11g — and a new 802.11n is on the horizon. Each of these technologies is based upon published standards, and the upcoming ones are usually based on several draft standards. Some hardware manufacturers have started to supply hardware that is capable of supporting the upcoming standards and using some of the new features. Typically, if the manufacturer can meet the hardware standards, then the software can be updated later via a firmware update.

To go along with the standard features, many vendors choose to implement additional features that are not in the standard or make propriety enhancements, like many of the vendors who advertise special speed enhancements.

The following sections outline the main features of each of the 802.11 wireless networking standards.

802.11a

Released in 1999 (but not actively shipping until component supply issues were resolved), 802.11a networks operated in the 5 GHz radio spectrum. Other devices, such as newer cordless phones, also run in this unlicensed spectrum. Overall, the 5 GHz space is less occupied by devices when compared to the 2.4 GHz space used by most other wireless devices. For 802.11a networks, there are 12 non-overlapping signal channels.

The network rated speed of 802.11a is 54 Mbps, but in most situations, you can expect to see about half that. The typical range for this type of network is 30 meters (nearly 100 feet).

The 5 GHz range may be unlicensed, but its use for wireless networking has been approved only in the United States, Canada, Japan, and prior to 2002, only in some European countries. Regulatory changes in 2002 and 2003 opened the 5 GHz range throughout the European Union.

802.11b

Released in 1999, 802.11b technology was able to beat 802.11a technology to the market because it was built on *DSSS (Direct-Sequence Spread Spectrum)* technology, and as such, components required for it were all readily available for production facilities. 802.11b operates in the 2.4 GHz radio spectrum, which is heavily cluttered with cordless phones, Bluetooth (see the “Bluetooth” section, later in this chapter), and spill-over interference from microwave ovens.

The network rated speed of 802.11b is 11 Mbps, but in most situations, you can expect to see about 6.5 Mbps because 802.11b is a *CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance)* network, which sends “I’m going to send data” broadcast messages, prior to sending its data, which increases the traffic on the network. (see Book VIII, Chapter 1, for the low-down on CSMA/CA). The typical range for this type of network is 30 meters (nearly 100 feet).

Wireless networks in the 2.4 GHz range have 14 channels available for use; 11 channels are acceptable everywhere, while the United Kingdom and some European countries allow up to 13 channels, and Japan allows all 14 channels to be used. *Channels* are defined by the center frequency that they use, and some vendors use wider channels than others, overlapping their neighboring channels. Typically 1, 6, 11, and 14 are thought to be non-overlapping channels, so in North America, that means that there are three channels that are normally non-overlapping and will not interfere with each other. See Figure 2-3 for a channel diagram of the 802.11b/g spectrum.

802.11g

Released in 2003, 802.11g was an updated version of 802.11b, still running in the 2.4 GHz spectrum but allowing faster data transfer speeds to put it on a performance par with 802.11a. Because it operates in the 2.4 GHz spectrum, it is also backward compatible with 802.11b, making migration to 802.11g an easy step for many organizations because they could update the back-end infrastructure and then update the wireless clients in a manner that fits within the corporate requirements.

Channel 1 2.401 – 2.423 Center 2.412	Channel 6 2.426 – 2.448 Center 2.437	Channel 11 2.451 – 2.473 Center 2.462	Channel 14 2.473 – 2.495 Center 2.484
Channel 2 2.404 – 2.428 Center 2.417	Channel 7 2.431 – 2.453 Center 2.442	Channel 12 2.456 – 2.478 Center 2.467	
Channel 3 2.411 – 2.433 Center 2.422	Channel 8 2.436 – 2.458 Center 2.447	Channel 13 2.461 – 2.483 Center 2.472	
Channel 4 2.416 – 2.438 Center 2.427	Channel 9 2.441 – 2.463 Center 2.452		
Channel 5 2.401 – 2.423 Center 2.412	Channel 10 2.446 – 2.468 Center 2.457		

Figure 2-3:
Normal
channel
frequencies
used for
802.11b/g
networks.

The network rated speed for 802.11g is 54 Mbps, but in most situations you can expect about half that, since the throughput is reduced as distances between the client and access point increase and the signal suffers from greater interference. The fact that wireless networks are a shared medium further reduces throughput as the number of users sharing an access point increases the compete for the wireless signal. The typical range for this type of network is about 30 meters (nearly 100 feet).

With all of the benefits that are available with 802.11g, many companies began producing products well before the standards were ratified.

802.11n

Due to be released in 2007, 802.11n represents the latest improvements to networking in the 2.4 GHz spectrum. Again, this solution will suffer from an already congested spectrum. The biggest change that you will see with this technology is the implementation of multiple transmitters and receivers and multiplexing of signals. This technology is referred to as *MIMO (Multiple-Input Multiple-Output)*.

The expected network-rated speed is 540 Mbps, but as with most wireless technology, you should expect typical transfer rates at about half that — in the 200 Mbps area. The typical range for this type of network will be about 50 meters (about 165 feet).

When a proposal request was announced for the 802.11n wireless networking standard, 32 responses were sent in. After voting by the standards committee, four of these proposals were considered further. In 2005, the three groups backing the three remaining proposals announced that they would work on a joint proposal to be accepted as the standard. Since there are multiple groups working together to design the final standard, there remains a possibility that the current draft standard may be substantially changed. Some vendors have started producing and shipping equipment based on the current draft standards.

Securing wireless networks

With the proliferation of wireless technology, security is starting to come to the forefront of many conversations, which is odd — it should have been prominent from the very beginning. The main steps you can take to improve wireless security include password management, SSID management, MAC address filtering, WEP, WPA, WPA2, DMZ, DHCP settings, and updating the firmware. Many of these configuration settings can be seen in Figure 2-4. The configuration screens for most access points are accessed using a Web browser and connecting to the IP address of the access point, as I did when connecting to the Linksys access point shown in Figure 2-4.

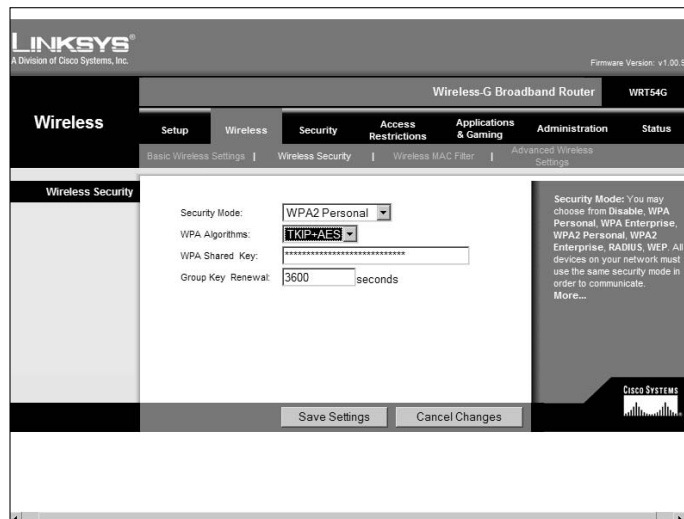


Figure 2-4: Security options for a typical wireless access point.

Passwords

All access points ship with a default Administrator username and password. A quick search of the Internet will give you the default usernames and passwords for most manufacturers and models. Usually, in addition to an owner's manual, a new access point will ship with a Quick Start guide which gives you an abridged set of steps that you need to follow to apply a basic configuration to the access point. This guide will typically include the default IP address, username, and password used by the access point. When wireless networks started to become widespread, these guides did not suggest changing the default passwords, but most of them now have the suggestion in the initial setup steps.



Changing the Administrator password should be the first step in the setup procedure. Until you change the password, any person within the range of your access point can make any configuration changes they wish.

SSID broadcasting

Security through obscurity is not the best security measure, but it provides one more layer to your overall network security. The Microsoft Wireless Zero Configuration service tries to make it easy for you to connect to wireless networks (as do many hardware vendors that have custom management software), and these configuration tools will display all of the SSIDs for wireless networks in range of your computer if the access points are broadcasting their SSID. This allows you to select the SSID or wireless network that you want to connect to and quickly configure it. By disabling the broadcasting of your SSID, your network will not show up in the list of detected wireless networks within Windows XP. If you are not on the list, most people won't even know your network is there and therefore won't connect to it.



Disabling SSID broadcasting is only a mild security setting, as you discover in the section, "Troubleshooting Issues," later in this chapter. If someone wants to gain access to your wireless network, disabling SSID broadcasts won't slow him or her very much or prevent that person from knowing that a wireless network exists.

MAC address filtering

MAC address filtering is another step that will slow most casual users from gaining access to your network, but it is not considered strong security on its own. Most access points allow you to perform different types of filtering, and most allow you to at least filter traffic by *MAC (Media Access Control) addresses* that are hard-coded onto each network adapter on your wireless network.

Implementing this security step means that for each new network card that you want to allow to use your network, you need to adjust your access points. Even with the MAC address restriction in place on your access point, by using a network sniffer like AirSnort, an unscrupulous person can see some traffic on the wireless network. In this traffic, such a person can see the MAC addresses of clients that are communicating on your wireless network. By taking this information and using a network card or operating system that allows him or her to manually control the MAC address, the individual can imitate one of the valid MAC addresses that are allowed to be used on your network. The timing on this is important because if two computers on a network have the same MAC address, then all sorts of communication errors can happen.



As mentioned previously, MAC address filtering keeps the casual user out of your network, but it shouldn't be your sole security measure. If somebody wants to get onto your network, this will not slow him or her down very much.

WEP

The previous security measures do not deal with encryption. *Wired Equivalent Privacy (WEP)* was the first attempt to secure wireless networks by using encryption. Early WEP used a 64-bit key, which was a 40-bit encryption key joined with a 24-bit *initialization vector (IV)*. This was easily breakable. Newer WEP uses a 104-bit encryption key joined with the 24-bit IV, providing a higher level of security. The benefit of WEP is that it's fairly easy to configure on both the access point and the client and again provides security against the casual wireless user. The main drawback is that technology has not substantially changed on how security keys are managed (such as IVs being sent over the network in plain text), and if an intruder captures enough traffic from your network, he can deduce your WEP keys, giving him full access to your network. The answer to this problem is WPA.

WPA

Wi-Fi Protected Access (WPA) and *WPA2* (also called *802.11i*) are all about improving security on wireless networks. Rather than using a single WEP key, WPA uses per-session keys that are agreed upon by the wireless client and the access point after an initial handshaking process. This technology was created in response to the weaknesses that were found in WEP. No security is unbreakable, and the creation of a better mousetrap just seems to create smarter mice. But the security settings available in WPA2 make wireless networks as secure as they can be, and currently, the time that it would take to break the keys used would render the broken keys mostly useless because they would have already been changed on the network.

WPA allows you to use two initial *handshaking methods*, which are either an 802.1x authentication (Radius) server, or a manually typed pre-shared key (PSK), which is configured on both the access point and the clients. In your list of security methods, you can choose *TKIP (Temporal Key Integrity Protocol)*, which uses a 128-bit encryption key and a 48-bit IV to secure the data. In addition, after every 10,000 packets of data have been sent, new TKIP keys are generated and used. This regular changing of the keys keeps the network more secure.

WPA2 builds upon this security by adding the U.S. government standard *AES (Advanced Encryption Standard)* to the data encryption methods, which allows for encryption keys of up to 256 bits.



Not all wireless products support WPA and WPA2, but the list is getting longer every day.



WEP and WPA are the two most common security methods used to secure wireless networks.

VPNs and DMZs

Although it requires more setup, you can treat your wireless network as a hostile remote network. Some companies will have faith that the security settings that they have configured on their wireless access points will keep all unwanted visitors or trespassers off of their network, and that is not unreasonable, given the current security settings that are available. For those companies that do not trust these security levels available in the current wireless technology, or who are afraid of the smarter mouse arriving on the scene sooner rather than later, there are other steps that can be taken to provide even more secure wireless.

Most companies today have a connection to the Internet coming into their organization, and they treat the Internet as a large hostile remote network. Hostile in the sense that it contains many individuals who would like to gain access to the company's private information that is stored on its corporate network. To prevent their private data from being accessed, they use a firewall at the point where the Internet meets their network, allowing them to inspect and control the data that moves between their network and the Internet. If they have data that they would like some people on the Internet to be able to access, such as data on a Web server or ftp server, then they may implement a second firewall making another private network, which provides some protection from unwanted visitors gaining access to the Web or ftp server, but offers less protection than their normal private network, since some people have a right to access the server. In most cases, this more open private net will be placed between the hostile Internet and the private corporate network, creating a less hostile, but still not absolutely safe *Demilitarized Zone* or *DMZ*.

Many companies have remote workers who need access to the corporate network while they are away from the office. This access is given through a special secure channel called a *Virtual Private Network* or *VPN*. A VPN relies on secure authentication and data encryption methods to create a secure and private connection through a firewall to a corporate network. Most often this security is handled by SSL (Secure Sockets Layer) or IP-Sec (IP-Security), which are two industry standard methods of encrypting TCP/IP or Internet data.



More information about firewalls, VPNs, DMZs, SSL, and IP-Sec can be found in Book IX, Chapters 1 and 2.

So, by segregating your wireless users into their own DMZ or private segment (protected from the Internet) you can keep any potential wireless intruder away from your private corporate data, since users of that wireless network do not have direct access to the corporate network. In order for your users to access corporate data, they would use the same remote access methods, such as a VPN connection, that they would use when they are on the road. These additional security steps can be layered on top the previously discussed techniques, such as WEP and WPA, forcing users to take more steps to get access to your corporate data, but making the connection more secure.

DHCP server settings

A *DHCP* (*Dynamic Host Configuration Protocol*) server provides automatic TCP/IP configuration to network clients by changing normal items that DHCP provides, such as a default gateway off of your network, or DNS settings for name resolution, or by disabling DHCP altogether. By failing to distribute accurate TCP/IP information to unwanted visitors, and using manual TCP/IP configurations to users, you are implementing another rudimentary security step. This security based on DHCP settings may prevent most casual users from getting ready access to information or gaining access to the Internet through your wireless network, but again, any unwanted intruder who wants to gain access will be able to find correct TCP/IP settings by using a standard packet capture utility, such as Ethereal (www.ethereal.com).

Firmware

Not to be forgotten when setting up new wireless networks (or periodically after a network is set up) is checking for firmware updates for your wireless network components or driver updates for your network cards. These updates are how new security features are retrofitted into your wireless products.



Some care should be taken when doing this because features are occasionally removed or there are problems with the new code in the firmware.

Other Wireless

Two more products fall into the wireless networking category: Bluetooth and cellular. Bluetooth is used as a communication link between different devices, while cellular is used for remote dial-up networking or Internet access. In both cases the use of these wireless devices is very different than the 802.11 wireless networking protocols, which are designed to be a short range extension of your LAN (Local Area Network).

Bluetooth

Bluetooth follows the standards set out in the IEEE 802.15.1 specification. It is a *Personal Area Network (PAN)* and is designed primarily for use in very small areas or short ranges, such as on a person's body. The most common item you may notice on the street is the cordless cellular phone headset, which operates over the space between the headset in their ear and the phone on their hip or in their purse. Many of the range wireless links for cell phones and PDAs (Personal Digital Assistant) are using Bluetooth. You may not immediately think of these links as being on a network, since you are not surfing the Internet or transferring files, but you are transferring data, in some form, between the devices. There are three classes of devices on Bluetooth networks:

- ◆ **Class 1** devices operate at 100 mW and have a range of 100 meters.
- ◆ **Class 2** devices operate at 2.5 mW and have a range of 10 meters.
- ◆ **Class 3** devices operate at 1 mW and have a range of about 1 meter.

Most devices are Class 3. Class 3 Bluetooth devices include telephone and computer headsets, PDA-to-PC synchronization, printers, digital cameras, cell phones for synchronizing with PCs, game controllers, remote controls, and instrument collection devices.

Bluetooth has replaced serial or infrared connections that would have been used for many of these functions in the past. Bluetooth is used mainly in low-bandwidth, short-haul situations.

Bluetooth operates in the 2.4 GHz spectrum, specifically around the 2.45 GHz mark. This means that it is open to interference from other devices running in the same spectrum, although due to its limited range, it tends to cause fewer problems with devices which are outside of its range.

Some Bluetooth devices have a security feature called *pairing* that can and should be used when available. There have already been viruses affecting Bluetooth cell phones which originally were set to freely associate with any other Bluetooth device in the area. Most phones now require pairing with devices to communicate. *Pairing* registers a pair of devices with each other by using a shared secret key, so that they will only be able to talk to

other devices which are known. Pairing is used as a means of authentication between devices and can also be used to encrypt data communication between them. Some devices, such as printers, may end up being left open and unsecured to keep a high level of functionality, but this would be a conscious decision to leave pairing turned off, and to leave the device unsecured.

WAN cellular

A *WAN* is a *Wide Area Network*, and unlike a LAN which is contained within one or two adjacent building, a WAN spans a large geographical area, and usually involves communication links that are operated by a Telco. Considering the number of communities and cities planning to set up wide-scale public access wireless networks and the cost of cellular data transfer rates, WAN cellular technology may not become extremely widespread. In conjunction with your data network provider, you can sometimes get cellular WAN access to your network which is also called Wireless Wide Area Network (WWAN). You need a special cellular gateway or a data enabled phone, which some people may refer to as a modem. Unlike a normal cellular modem which needs to dial another modem to establish a connection, the gateway makes a connection directly to the cellular provider's network. The gateway may be connected to your computer by USB, Bluetooth, or PCMCIA. This gateway connects to your data network provider and sets up a secure tunnel for access to your corporate network's resources. The data rates on these connections are usually faster than traditional dial-up connections and they provide secure access to your LAN data services, but most service providers charge a substantial fee for the service.

Increasing Wireless Network Performance

Some basic steps that you can take to improve wireless performance include the following:

- ◆ **Reduce interference from other sources.** This might include changing channels that you are using or changing the spectrum that you are using, such as switching from 2.4 GHz to 5 GHz technologies.
- ◆ **Reduce the number of clients per access point.** Because wireless is a shared medium, the fewer clients the access point has to share bandwidth between, the faster each client will be.
- ◆ **Reduce the range between the clients and the access points.** As the range between the client and access point increases, the signal quality decreases, and therefore the data throughput rate decreases.
- ◆ **Use faster technology.** For example, upgrade from 802.11b to 802.11g or 802.11n.



- ◆ **Reduce the material that the signal must travel through.** The distance ratings are usually open-air ratings, and construction materials and other objects in the path of the signal will reduce the signal strength.

Outdoor wireless networks set up during the winter can experience problems in the spring when all the trees in the path of the signal now have leaves on them. Try to anticipate how annual weather and seasonal changes in your area will affect both your wireless equipment and wireless signal.

- ◆ **Use proprietary or non-standard technology to improve data transfer speeds.** This may work only when teamed with specific network cards from the wireless vendor.

Troubleshooting Issues

Any issues that affect wired networks can affect wireless networks as well, and very often, symptoms show up first on the wireless network. The bandwidth on the wireless network is substantially lower, so when there is a problem on your network that consumes bandwidth, like a *worm*, it will cause connectivity and performance issues on your wireless network before users of your wired network notice the problem.

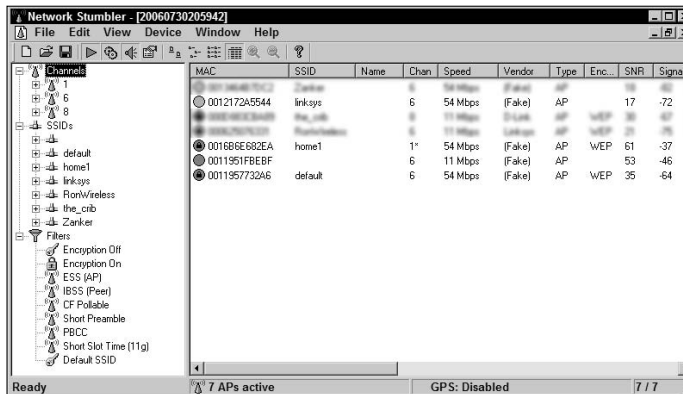
A number of outside factors, such as microwaves and cordless phones, can affect wireless networks. Since these devices operate throughout and beyond the 2.4 GHz spectrum, and they do not advertise themselves as wireless networks, they won't show up in the results in Network Stumbler when you're viewing the strength of wireless networks in the area. In order to locate these rogue signals, you need to either perform a physical inspection or, better yet, run a *spectrum analysis tool*. In the past, these tools cost thousands of dollars, but recently, quality tools have dropped into the hundreds of dollars, and some tools are even less than a hundred dollars. Managing your wireless frequencies is becoming just as important as managing your wired data networks.

If you want to survey wireless networks in your area so that you can choose a clear channel, take a look at Network Stumbler (www.netstumbler.com). It lets you view detailed information about the wireless signals in your area. Figure 2-5 shows how Network Stumbler displays information about access points, including the channels on which they are running, what SSIDs are broadcast (AP 0011951FBEBF is not broadcasting its SSID), and whether they have security enabled. Many other applications do the same type of thing, and this feature is even incorporated into the management software that ships with some network card drivers.

If you know that there is an access point that is not broadcasting its SSID, you can find the SSID by using wireless network auditing tools, like the BackTrack (formerly Linux Auditor Security), which is available on the www.remote-exploit.org Web site.



Figure 2-5:
Network
Stumbler
and similar
tools let you
see wireless
networks in
your area.



If you can get a good signal but you don't seem to be able to establish a connection or transfer data with the network, then you should look at the security settings on both the access point and the client computer. If the WEP settings are incorrect or if the network authentication is not set correctly, then you may experience either of these problems.

Much the same as WEP keys and settings, if you are using WPA, your problems could be with the shared secret that is configured or the settings for the authentication server, which verify user or access point credentials.

Getting an A+

This chapter goes over some common wireless communications. The following points are covered:

- ◆ Infrared technology depends on line of sight and tends to be used for low-bandwidth solutions.
- ◆ A number of factors affect wireless signal quality and strength, such as interference, range, and the number of connected clients.
- ◆ 802.11a/b/g networks all have a range of 30 meters or 100 feet and while 802.11a operates in the 5 GHz spectrum, 802.11b and 802.11g operate in the 2.4 GHz spectrum.
- ◆ WEP and WPA are the most common methods to secure wireless networks, but other options include disabling SSID broadcasting, MAC filtering, placing wireless APs in a DMZ, and changing default passwords.
- ◆ Bluetooth is used only for short range communication.

Prep Test

1 Which of the following is not a technology that can be used to connect to a printer?

- A Bluetooth
- B WWAN
- C 802.11b
- D Infrared

2 Typical client Bluetooth implementations have what ranges? (Select all that apply.)

- A 1 meter
- B 5 meters
- C 10 meters
- D 25 meters

3 What is the standard range of 802.11a/b/g wireless networks?

- A 10 meters
- B 30 meters
- C 50 meters
- D 100 meters

4 Which of the following does *not* impact the quality of wireless signals?

- A Microwave ovens
- B Air temperature
- C Cordless phones
- D Distance between access point and clients

5 802.11g networks are backward compatible with which other 802.11 technology?

- A 802.11a
- B 802.11b
- C 802.11i
- D 802.11n

Common Wireless Communications

- 6** Which of the following would be used to provide wireless connectivity across a metropolitan area?
- A Bluetooth
 - B Infrared
 - C Cellular
 - D MIMO
- 7** What frequency ranges are used by 802.11 wireless networks? (Select all that apply.)
- A 900 MHz
 - B 1.5 GHz
 - C 2.4 GHz
 - D 5.0 GHz
- 8** Which of the following is not a technology used to secure wireless networks?
- A WPS
 - B WEP
 - C WPA
 - D 802.11i

Answers

- 1** B. Printers from different manufacturers support some or all of the listed wireless technologies, while WWAN is a technology used to connect to your network while out of the office. *See "Understanding Infrared Devices," "Bluetooth," and "Working with Wireless Networks."*
- 2** A, C. Bluetooth technology has ranges of 1, 10, and 100 meters depending on the class of devices that are being used. *Review "Bluetooth."*
- 3** B. Most 802.11 wireless networks have a range of 30 meters in open air. *Check out "Wireless standards."*
- 4** B. Air temperature has not been documented to have an impact on wireless networks. *Peruse "Troubleshooting Issues."*
- 5** B. 802.11g uses the same 2.4 GHz spectrum and was designed to be backward compatible with 802.11b. 802.11n is expected to be backward compatible with both 802.11g and 802.11b. *Take a look at "802.11g."*
- 6** C. The only technology that will work across a metropolitan area is cellular, which is used for WAN cellular or WWAN connectivity. *Peek at "WAN cellular."*
- 7** C, D. 802.11 networks run at either 2.4 GHz for 802.11b/g/n or 5.0 GHz for 802.11a. *Look over "Wireless standards."*
- 8** A. The technologies used for wireless security are WEP, WPA, and WPA2, which is also called 802.11i. *Study "Securing wireless networks."*

Chapter 3: Networking the Operating System

Exam Objectives

- ✓ Understanding networking components
- ✓ Identifying network protocols
- ✓ Configuring the TCP/IP protocol
- ✓ Troubleshooting with TCP/IP utilities
- ✓ Understanding name resolution
- ✓ Sharing network resources
- ✓ Connecting to shared resources

Today, one of the most important troubleshooting skills IT professionals can adopt is the ability to troubleshoot networking connectivity. *Network connectivity* is the term used for two computers establishing a connection to one another. A number of components allow this communication to happen, and a lot of times, it is these components that IT professionals end up troubleshooting.

This chapter discusses the software components that allow a computer to network with other computers and the troubleshooting issues that could arise while configuring the network. In this chapter, you find out how to connect to networking resources and how to troubleshoot when you can't make the connection.

Understanding Networking Components

When setting up a network, you must have the appropriate hardware and software in place to allow systems to communicate with one another. Since this chapter focuses on the software components that are needed to allow Windows to network, you can assume that you have all of the necessary hardware in place. You have purchased a hub or switch, you have at least two computers and network cards to go in the computers, and the appropriate cabling to connect the network cards to the hub/switch are already connected. After all the hardware is in place, what do you have to do at the operating system level to get these computers talking? The answer — not only do you need the physical hardware in place, but you also need to load

software components such as a network card driver, protocol, service, and client software. These are the four major software components required to network.

When building your network, it is important to identify the four major software components that allow a Windows operating system to function in a networking environment. These components are

- ◆ **Network adapter driver**
- ◆ **Network client**
- ◆ **Protocol**
- ◆ **Services**

I discuss each of these components in detail in the sections that follow. Be sure to be comfortable with them for the A+ exams!

Network adapter driver

The *network adapter driver* is the physical network card that is inserted into one of the computer's expansion bus slots, connected as a USB device, or integrated into the system board. The network card is responsible for sending information out onto the network and receiving information from the network.


Before purchasing a network card, you have to figure out what *type* of card you need. To do this, you need to open up the computer (for more information on safety procedures, refer to Book I, Chapter 3), look at the expansion buses that are supported in your system, and then identify which has an empty slot. For example, you may open your computer and see that you have an ISA slot and three PCI slots, so you have a choice of purchasing a PCI or an ISA network adapter. Typically, you would purchase the PCI network adapter because of the performance benefits of PCI devices over ISA devices, along with the fact that you will be hard pressed to find an ISA card at the store nowadays.



A USB network adapter is a popular choice today. If you purchase a USB network adapter, you won't need to open the computer — simply plug the USB adapter into an available USB port.

After you insert the network card into the empty expansion slot or USB port, you need to install the driver for that card within the operating system. Installing the network card driver is the first major step to networking a system. The *driver* software allows the operating system to communicate with the physical device, which in this case is the network card. Figure 3-1 shows how the driver sits between the operating system and the physical hardware, controlling communication between the two.

Figure 3-1:
The relationship between the operating system and a hardware device is controlled by the device driver.



Installing a network adapter in Windows 2000/XP/2003

After you have inserted the network card into the computer, you will probably notice that Plug and Play will kick in when the computer is turned on. If the operating system has the device driver for the card, it will load the driver up automatically and you will see the device listed in Device Manager.

If the operating system doesn't have the driver, it will either prompt you for the driver (and you will need to supply the manufacturer's CD for the network card) or the device will be listed in Device Manager as an unknown device. To update the driver for the device in Windows 2000/XP/2003, you need to go to Device Manager to update the driver. The following steps demonstrate how to update a driver in Windows XP:

- 1. Click Start, right-click My Computer, and choose Properties.**
- 2. In the System Properties dialog box, click the Hardware tab.**
- 3. On the Hardware tab, click the Device Manager button.**

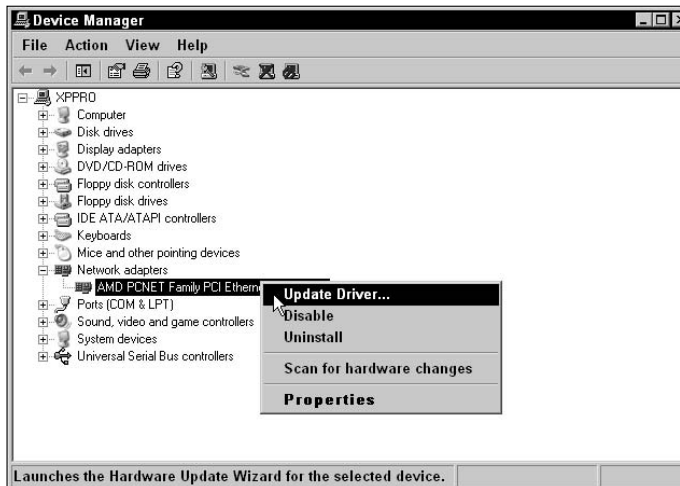
The Device Manager appears.

- 4. Within Device Manager, right-click on your network card and choose Update Driver, as shown in Figure 3-2.**

If your network card isn't listed in the Network Adapters section of Device Manager, look under the Unknown Devices category. If the device is in the Unknown Devices category, right-click it there and choose Update Driver.

The Update Driver Wizard starts.

Figure 3-2:
Updating a
network
card driver
in Windows
XP.



5. Select **Install from a List** and click **Next**.
6. Select **Don't Search, I Will Choose the Driver to Install**; then click **Next**.
7. Select the manufacturer of your network card on the left and then choose the model of your network card on the right. Click **Next**.

If your network card model isn't in the list, click the **Have Disk** button so that you can supply the location of the driver.

8. After the driver is copied, click **Finish**.

Understanding the local area connection

After you load a network card driver, Windows creates an icon that represents the network card; this icon is called the *local area connection*. If you have multiple network cards installed, you will have multiple local area connections — one representing each network card.

The purpose of the local area connection icon is to give you a place to configure any network settings that are responsible for communication between the network card and the local area network (LAN). For example, if you want to ensure that TCP/IP is being used by your network card, you right-click on the local area connection and choose **Properties**. Once in the properties of the local area connection you can add or configure networking components such as TCP/IP.

To view your LAN connections in Windows XP, choose Start⇨Control Panel⇨Network and Internet Connections⇨Network Connections. You will see a window that displays your local area connection icon, as shown in Figure 3-3.

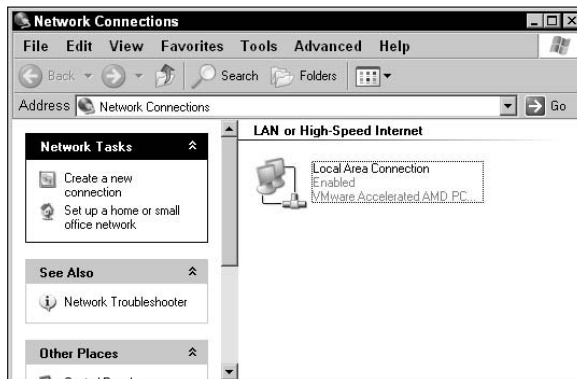


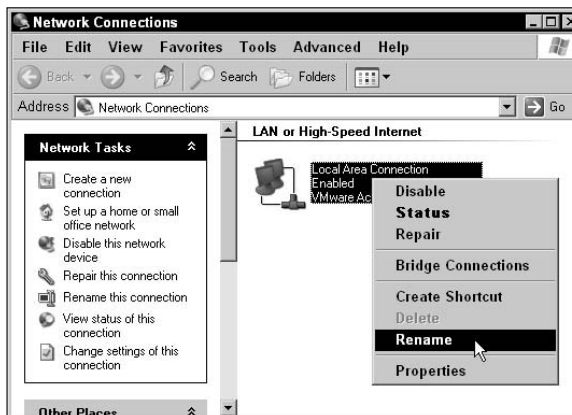
Figure 3-3:
Viewing a
local area
connection
in Windows
XP.

If you right-click the local area connection icon, a context menu gives you a number of tasks. The following is a list of the tasks that you might use when you troubleshoot networking issues:

- ◆ **Disable:** Choose Disable if you want to temporarily cut off communication to and from the network. This is a quicker solution than physically removing the network card from the computer.
- ◆ **Enable:** When you disable a card, the Disable option changes to Enable. After you troubleshoot your network (hopefully, you fix the problem) and you would like to re-enable the network connection, choose the Enable command.
- ◆ **Status:** The status command displays a dialog box that shows how long the connection has been up and running and also the speed of the connection. On the Support Page tab, you can view your IP address information and MAC address.
- ◆ **Repair:** If you click the Repair command, Windows performs maintenance on the connection by performing tasks such as renewing your IP address, flushing the ARP cache, and flushing the NetBIOS and DNS resolver cache.
- ◆ **Rename:** Use this command to give the connection a more meaningful name. For example, I renamed the local area connection for the network card that is connected to the Internet “Internet Connection.” The Rename command is shown in Figure 3-4.

- ◆ **Properties:** Use the Properties command to open the Properties dialog box for your LAN connection. In the Properties dialog box, you can modify the network setup of the network card. For example, you can add or remove network protocols or change their configuration.

Figure 3-4: Renaming a local area connection to give it a more meaningful name.



After you make sure that the correct hardware settings are applied to the network adapter, your next step is to connect to a network resource. Unfortunately, you can't connect to a network resource until you have the appropriate network client running. The following section describes the purpose of the network client.

Network client

A *network client* is no different than a client or customer in the real world. A client in the real world visits your company because you provide some sort of service. For example, you might run a tailor shop, which provides a particular service to customers who drop off pants or dresses that need tailoring. In a sense, the customer is a *client* of the tailoring service.

Computer networks work the same way. On your computers, you must run a client for the type of service you are requesting on the network. For example, if you work for a medium-sized company that runs Novell's NetWare as the server operating system, then you must load a client that will connect your computer to the Novell server. Or, if you want to connect to a Windows 2003 server, you have to load a Microsoft client on your system to do so.

The Windows operating systems come with two major clients already installed: Client for Microsoft Networks and Client Service for NetWare. When you want to connect to a Microsoft server, you need to have the Client

for Microsoft Networks software loaded, and when you are connecting to a Novell server, you load the Client Service for NetWare (CSNW).



If you are running Windows operating systems in a Novell environment, you will probably decide *not* to load the Client Service for NetWare, the client that Microsoft has built in to its operating system to connect to Novell networks. Microsoft has built that client with limited capabilities, so most networks that run Novell actually load Novell's Client on the Windows desktops instead.

To return to the tailor shop example, remember that your client has asked your service to tailor some pants. When the client finally receives the mended pants, that client is pleased — however, the pants have to be dry-cleaned before they are used. Unfortunately, your business doesn't offer dry-cleaning services, so your client has to request the service from a third party. The point is that your client can be a client of tailoring and a client of dry cleaning at the same time. There is no rule that says you can be a client of only one particular service at a time.

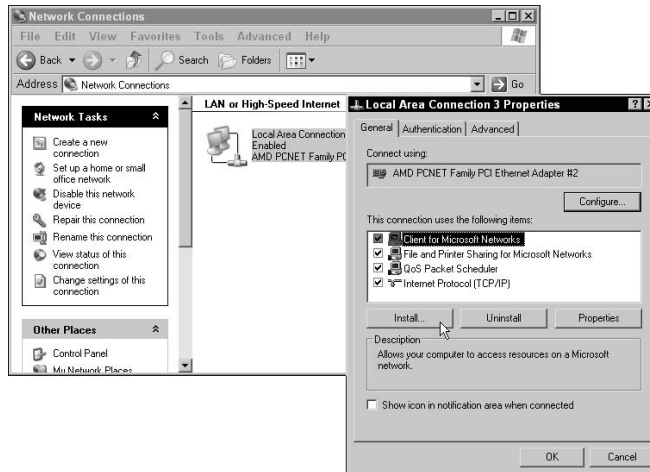
This applies to the network as well. A lot of companies run both Windows servers and Novell servers on the same network and at the same time. Maybe they use the Novell server for file sharing and use the Windows server for e-mail services. In this instance, the desktop computers on the network would have to run two clients, Client Service for NetWare — or the Novell client software, and Client for Microsoft Networks.

To install a network client for Microsoft networks in Windows 2000/XP/2003, follow these steps:

- 1. If you're using Windows 2000, choose Start⇨Settings⇨Control Panel⇨Network and Dial Up Connections. If you're using Windows XP/2003, choose Start⇨Control Panel⇨Network and Internet Connections⇨Network Connections.**
- 2. Right-click your local area connection and choose Properties.**
- 3. If it isn't already selected, click the General tab.**
- 4. Verify that Client for Microsoft Networks appears in the list of components and has a check mark beside it, as shown in Figure 3-5.**

If the check box is not selected, click the check box beside Client for Microsoft Networks to enable the Microsoft networking client. If the Client for Microsoft Networks is not in the list, install it by clicking the Install button and choose Client from the Network Component Type dialog box. After choosing client and clicking Add, you choose Client For Microsoft Networks and then click OK.

Figure 3-5:
Verifying
that
Client for
Microsoft
Networks is
installed on
a Windows
XP system.



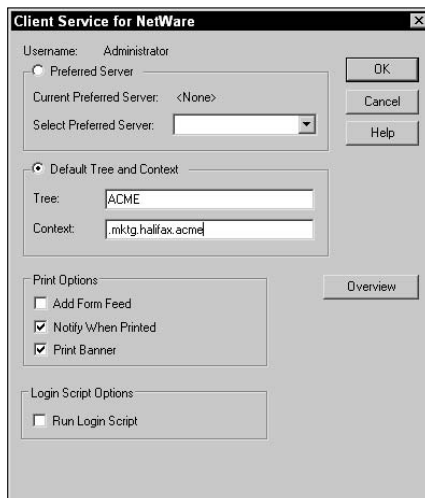
The steps are very similar to install Client Service for Netware and to have a Windows system connect to the Novell server on the network. To install Client Service for Netware, follow these steps:

1. If you're using Windows 2000, choose Start→Settings→Control Panel→Network and Dial Up Connections. If you're using Windows XP/2003, choose Start→Control Panel→Network and Internet Connections→Network Connections.
2. Right-click your local area connection and choose Properties.
3. Click the Install button to install a network component.
4. From the Network Component Type dialog box, select Client and then click Add.
5. Select Client Service for Netware and then click OK.
6. Click Yes to restart the system.

When you reboot after installing Client Service for Netware, you need to configure the client for your Netware tree and the context within that tree. The Client Service for Netware dialog box, shown in Figure 3-6, appears on reboot or can be configured later through the Client Service for Netware applet in the Control Panel.

The *tree* setting is the name of the Netware tree that holds your user accounts, and the *context* setting is the Organizational Unit (OU) in that tree that holds your user account. *Organizational units* are containers that hold objects in the directory.

Figure 3-6:
Configuring Client Service for Netware to authenticate with a particular Netware tree and context within that tree.



At this point, you should have your network card and its driver installed, and you should also have installed the appropriate client. Unfortunately, you are still unable to communicate with someone on the network because you still have to install the appropriate *protocol*. In the following section, I talk about the purpose of protocols and what common protocols are running on networks today.

Protocols

Protocols are languages that are used to hold a conversation on the network. Your system can have a network card installed and have the proper client running, but if it isn't speaking the same language (protocol) as the remote system, then the two systems can't hold a conversation. To go back to the tailor shop example, you are now ready to service your clients. There is only one problem: When your first client walks into your store and requests service, your client speaks French, while you speak only English. To solve this problem, you and your client must speak a common language. It doesn't matter what that language is, as long as you both can speak it.

There are a few things to look for when choosing which protocol to install, but the bottom line is that all computers on the network must have the same protocol installed — a common denominator to allow all individuals to participate in a conversation. In the following sections, you can examine a few different protocols that you might encounter when working with networks. Be sure to be familiar with these protocols when you take the A+ Certification exams.

NetBEUI

NetBIOS Extended User Interface (NetBEUI) was originally developed by IBM to be used on small networks — less than 10 computers. Microsoft implemented NetBEUI in the different Windows operating systems for the same purpose — small networks. NetBEUI is intended for small networks because it is a nonroutable protocol, meaning that it cannot leave the network. Since many companies have large networks spanning some form of WAN (Wide Area Network) link and containing routers to connect different networks, this protocol is impractical in those environments — and for environments that want to connect to the Internet. In today's networking environments it is unlikely to see NetBEUI on any major network.

In the past, you would use NetBEUI if you had a small number of computers that needed to be networked in a workgroup-type environment. You wanted to get this network up and running without the hassle of having to configure all kinds of settings. This is the benefit of NetBEUI: There is no configuration — it just works!

IPX/SPX

Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) is Novell's native protocol and has been used on large networks since its design. IPX/SPX is a *routable protocol*, meaning that the data it delivers can move from one network to another. Because the Internet is so popular, even Novell networks today don't run IPX/SPX but run TCP/IP instead (see the following section for information about TCP/IP).

Microsoft operating systems have their own implementation of IPX/SPX, known as *NWLink*, and you will want to install it if your computers communicate with Novell servers on the network.

Like NetBEUI, there is not a lot of configuration to IPX/SPX, but one very important property of this protocol is the frame type. The *frame type* is the type of "envelope" that is used to deliver the information from one computer to another computer on the network. Different envelopes have different characteristics, and you need to make sure that everyone is using the same type of envelope on the network — meaning that if you are running IPX/SPX, you need to ensure that you are running the same frame type on the client as on the Novell servers.

Windows operating systems can autodetect the frame type that is running on the network, which is usually first set by the Novell server. Although autodetect generally works out fine, it is important that you know how to install IPX/SPX and configure the frame type. The big question is, what frame

type do you use? The answer is probably whatever is set on the Novell server, but you should ask the network administrator. Popular frame types are 802.2, 802.3, and Ethernet II. You will not need to worry about the frame types for the exam, but be aware that the frame type must match that of the Novell server. For example, if the server is using 802.3 then the client computers must be configured for 802.3.

The following steps walk you through installing NWLink (IPX/SPX) on Windows 2000 and Windows XP and demonstrate how to set the frame type:

1. If you're using Windows 2000, choose Start⇨Settings⇨Control Panel⇨Network and Dial Up Connections. If you're using Windows XP/2003, choose Start⇨Control Panel⇨Network and Internet Connections⇨Network Connections.
2. Right-click your local area connection and choose Properties.
3. Click the Install button.
4. Select Protocol in the network component list and then click Add.
5. Select NWLink and click OK, as shown in Figure 3-7.

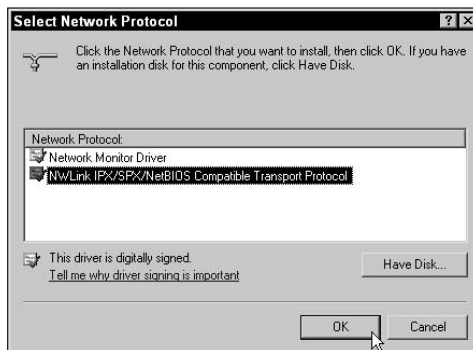


Figure 3-7:
Installing
NWLink on
Windows
XP.

6. When the NWLink protocol is in the list of network components used by the connection, select the NWLink IPX/SPX protocol (*not* NWLink NetBIOS) and click Properties to change the frame type.
7. Change the frame type setting from Auto to the value used by your IPX network.

For example, if you know that your Novell servers use 802.2, select 802.2 from the list. Check with your network administrator to find out what frame type your Novell servers use.

TCP/IP

Transmission Control Protocol/Internet Protocol (TCP/IP) is the hot protocol on the market these days because it is the protocol of the Internet and Internet-based technologies. TCP/IP has become the protocol of choice for Windows, Linux, and Novell networks because of its ability to communicate in heterogeneous environments. The bottom line is this: It doesn't matter what kind of operating system you are running — if you're running TCP/IP, you have the ability to communicate globally.

TCP/IP is known as a protocol suite. *Protocol suites* are like application suites in the sense that there is more than one protocol in the group. For example, if you purchase the Microsoft Office suite, you purchase an entire group of applications, or an entire *suite* of applications. TCP/IP is a group of protocols that make up the protocol suite, and some of these protocols are used day in and day out. For example, adding the TCP/IP protocol suite to your computer means that you have a telnet application for running applications from another computer, and you have an FTP application for downloading files from another computer.

Because configuring TCP/IP is such a big topic, I devote an entire section of this chapter, "The TCP/IP Protocol," to it.

For the exam remember that TCP/IP and IPX/SPX are routable protocols while NetBEUI is a non-routable protocol.



Services

One of the most forgotten networking components is the service. A *service* is a piece of software running on the computer that provides certain functionality. An example of a service that runs on the computer is file and printer sharing — which is the service that provides files, folders, and printers to other systems on the network.

Going back to the tailor shop example, before you can have clients, you must have first made the decision to offer the service. If you don't offer the service, then there would be no reason for customers to want to communicate with you.

On the network, someone has to offer the service, but not everyone needs to. For example, a small company with five Windows machines may have only the machine with the printer connected to it providing the file and printer sharing service. The other four Windows computers connect to it by installing Client for Microsoft Networks and ensuring that they are using the same protocol — there is no reason for them to have the service installed because they are offering nothing to the network.

Real-world scenario: Planning network components

You are responsible for determining what networking components are required in order to build the scenario below. The answer follows the scenario.

Scenario: You are in charge of implementing the network infrastructure for your organization. There are two Windows 2003 servers and three Novell servers servicing 200 Windows XP clients. The word processing department will be sending print jobs to one Windows XP computer in the manager's office. What networking components would you load to allow the clients to access both the Novell servers and the Windows 2003 servers?

Answer: The Windows XP computer in the manager's office will have File and Print Sharing for Microsoft Networks loaded so that the Windows XP client in the word processing department may connect to it.

You will have to load a Novell client so that the Windows XP clients can connect to the Novell servers. If the Novell servers are using IPX/SPX, you will also have to make sure that the IPX/SPX protocol is loaded and the proper frame type is configured.

To allow the Windows XP clients to connect to the Windows 2003 servers, you will have to load the Client for Microsoft Networks if it isn't already loaded.

Windows and Novell servers usually run at least two services by default: File sharing services and printer sharing services. *File sharing services* allow the server to share files with other users on the network. *Printer sharing services* allow a printer to be used by multiple users on the network — you don't need to purchase a separate printer for each user on the network, which is a real cost cutter!

File and printer sharing services were the original purpose in life for servers and networks, but the number of services that can be added to these systems has grown over the years to include mail services, Web services, FTP services, name resolution services, and many more.

The TCP/IP Protocol

Since the dramatic growth of the Internet, *Transmission Control Protocol/Internet Protocol (TCP/IP)* has become the preferred protocol on networks today. TCP/IP is the common protocol on all desktops, including Windows, Linux, and Macintosh systems — allowing all of these different operating systems to communicate over a common protocol. It doesn't matter what operating system you are running or what kind of network you have — as long as you are running a common protocol such as TCP/IP, you can access resources across any platform.

TCP/IP is installed by default with all the major operating systems such as Windows 2000, Windows XP, and Linux. When configuring TCP/IP on these systems there are three major settings that need to be configured to allow the computer to communicate with other computers on different networks or the Internet:

- ◆ **IP address**
- ◆ **Subnet mask**
- ◆ **Default gateway**

In order to troubleshoot communication across TCP/IP, it is important that you understand the types of settings that need to be configured. In the sections that follow, you look at how to configure TCP/IP and at some utilities to help you troubleshoot the protocol.

IP address

The *IP address* is a 32-bit number that is unique to your computer. No two systems can have the same IP address. An IP address is similar to the address of your home, which is the method by which other people send mail to you. An IP address works the same way on a TCP/IP network — you will assign the number to your computer, and it is the method other computers use to send information to your computer.

An IP address is made up of four sets of numbers separated by periods. This is called the *dotted decimal notation format* of the IP address. An example of an IP address in the dotted decimal notation format is 131.107.2.200. Each of the four sets of numbers is referred to as an *octet* (because each octet represents 8 bits of data).

The IP address is made up of a network ID and a host ID:

- ◆ The **network ID** is a unique number used only by your network and is the same for all computers on the same network.

For example, in the IP address 131.107.2.200, the first two octets (131.107) make up the network ID. So if Computer A with an IP address of 131.107.2.200 talks to Computer B, which has the IP address of 131.107.3.5, you can assume that the two computers are on the same network because 131.107.x.y is the network ID for both computers.

- ◆ The **host ID** portion uniquely identifies a computer on the network.


For example, in the IP address 131.107.2.200, the last two octets (2.200) make up the host ID. Only one computer on the 131.107.x.y network can have the host ID of 2.200.



Is the network ID always the first two octets? The answer is no. The network ID isn't always the first two octets of the IP address. So how do you know which octets make up the network ID and which octets make up the host ID portion of the IP address? What *class* of IP address you have dictates which numbers correspond to which IDs.

There are three major classes of IP addresses: *Class A*, *Class B*, and *Class C*. The different IP address classes support a different total number of workstations on the network. For example, a Class A network (a network using Class A addresses) supports up to 16,777,214 network devices, while a Class B network supports 65,534 network devices, and a Class C network supports only 254 network devices.

Class A networks use the first octet as their network ID and the remaining three octets as the host ID. A Class B network uses the first two octets for the network ID and the last two octets as the host ID. A Class C network uses the first three octets as the network ID and the last octet for the host ID.



How do you know what class IP address you have? Look at the first octet. If it has a value between 1 and 126, it's a Class A IP address; if it has a value between 128 and 191, it's a Class B IP address; if it has a value between 192 and 223, it's a Class C IP address. Table 3-1 summarizes IP address classes.

Network Class	Number of Hosts	Octet Summary	First Octet Value
Class A	16,777,214	n.h.h.h	1–126
Class B	65,534	n.n.h.h	128–191
Class C	254	n.n.n.h	192–223



You will notice that in the number of hosts column there are actually two numbers missing. For example, a class C address can have 256 possible addresses numbered 0 through 255. But you are not allowed to use the 0 because it is reserved for what is known as the network ID. Also, you are not allowed to use the 255 because it is the broadcast address — which is how systems send data to every computer on the network. To summarize, with each of the address classes you will lose two addresses due to the network ID and the broadcast address.

Notice in Table 3-1 that the number 127 is skipped in the First Octet Value column. This is because it is illegal for a system on the network to have an IP address that starts with 127. Any address starting with 127 is illegal because this address is reserved for the *loopback address*. The loopback address is an address that always refers to the TCP/IP software stack that initialized on *your* system. Typically, the loopback address is referred to as 127.0.0.1, but

you could use any address that starts with 127. For example, you could issue the following command in a command prompt to test that TCP/IP is functioning properly on your system:

```
Ping 127.0.0.1
```

No matter what your IP address is, the loopback address is a constant that you can use to verify that your system is running correctly.

Here's an IP address example: Computer A has an IP address of 194.12.11.10, and it initiates communication with Computer B, whose IP address is 194.13.11.9. Are the two workstations on the same network? Looking at the first octet, you can see that the number 194 is a Class C address, which means that the first three octets make up the network ID, and the last octet is the host ID. Since there is a difference in the first three octets (which is the network ID), these two computers are on different networks.



Identifying whether the computer you are trying to communicate with is on your network could be important when troubleshooting communications. If the computer that you are trying to communicate with is not on your network, then the problem could be with your computer, the remote computer, or the router. This means you have more places to look for the cause of the problem.

Subnet mask

Another way to tell whether your computer is on the same network as the computer you are trying to communicate with is to look at the subnet mask. The *subnet mask* is what your computer uses to determine whether the network device it is trying to communicate with is on the same network or not. The subnet mask helps the system determine the network ID portion of the IP address by comparing the subnet mask against the IP address. If there is a 255 in the subnet mask, then the corresponding octet in the IP address is part of the network ID. Once the network ID is known, any systems that have the same network ID are considered to be on the same network — otherwise they would have a different network ID.

Different classes of IP addresses are associated with different default subnet masks. For example, a Class A address has a default subnet mask of 255.0.0.0, a Class B address has a default subnet mask of 255.255.0.0, and a Class C address has a default subnet mask of 255.255.255.0. Looking at the subnet masks, any octet that has the value of 255 means that the corresponding octet in the IP address is part of the network ID.

To put this all together, Table 3-2 shows an example of two computers and their IP address configurations.

Table 3-2 Comparing IP Addresses with Subnet Mask

<i>Item</i>	<i>Address</i>
Computer A (IP Address)	13.10.12.120
Computer A (Subnet Mask)	255.0.0.0
Computer B (IP Address)	18.23.48.119

Using this example, Computer A tries to connect to Computer B. The first thing that Computer A does is compare its IP address with its own subnet mask to determine what octets make up the network ID. Here it sees that the first octet is the network ID because the subnet mask has the number 255 in only the first octet. Then Computer A compares its subnet mask with the IP address of Computer B (the remote computer it is trying to communicate with) and identifies that the network IDs of Computer A and Computer B are different — the two computers are on different networks.

When Computer A realizes that the remote computer it is trying to communicate with (Computer B) is on a different network, it starts to panic because it doesn't have the capability to send the information over to the other computer. Computers can pass information directly to other computers only if both systems are on the same network. So what happens? This is where the default gateway fits into the story.

Default gateway

When information has to be forwarded from a computer on one network to a computer on another network, a special network device called a *router* must be used. The router has a table that lists all the networks it knows about and the network ID associated with each of those networks. When the router receives information destined for a particular IP address, it checks its table of network IDs for a match. If a match is found, it delivers the information to the appropriate network.

How does the information get to the router so that it can be forwarded? Looking back to the example from Table 3-2 in the previous section, Computer A has information for Computer B, and Computer A realizes that Computer B sits on a different network. At this point, Computer A looks at its *default gateway*, which is the address of the router that will forward the information on to Computer B's network. The default gateway is a TCP/IP option configured on each workstation. Typically, all computers on the same network point to the same router.

Now that you are comfortable with the concepts of an IP address, subnet mask, and default gateway, you're ready to configure these options on a Windows operating system, which is covered in the next section.

Configuring TCP/IP in Windows 2000/XP/2003

To configure TCP/IP on a newer system, such as Windows XP, you need to go to your local area connection properties and configure the TCP/IP protocol. Be aware that all Microsoft operating systems today have TCP/IP installed by default. You simply need to configure the IP address on the system. To configure TCP/IP on a Windows 2000/XP/2003 system, follow these steps:

1. If you're using Windows 2000, choose Start⇨Settings⇨Control Panel⇨Network and Dial Up Connections. If you're using Windows XP/2003, choose Start⇨Control Panel⇨Network and Internet Connections⇨Network Connections.
2. Right-click your local area connection and choose Properties.
3. In the list of items used by the connection, select TCP/IP and click Properties.
4. To assign a static address, select Use the Following IP Address (as shown in Figure 3-8) and then type your computer's IP address, subnet mask, and default gateway in the corresponding text boxes.

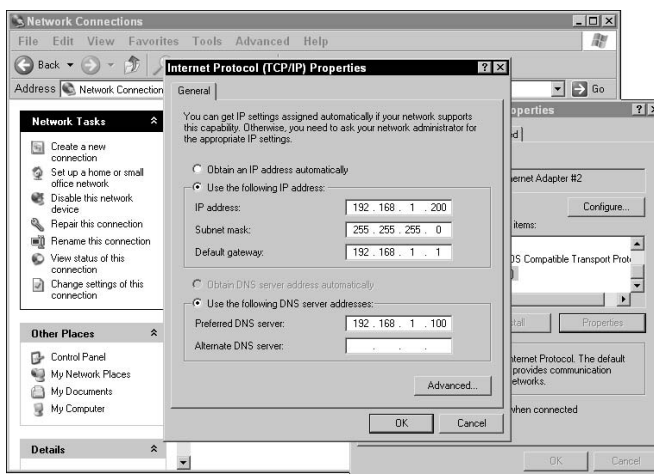


Figure 3-8:
Configuring
TCP/IP on a
Windows
XP client.

5. Type the address of your DNS server in the Preferred DNS Server text box.

In order to know what to type as the IP address of your DNS server you will need to consult the network administrator or maybe even the network architects. Whoever has designed the network knows the IP address of the DNS server. (DNS is covered in more detail in the section “DNS,” later in this chapter.)

6. Click OK and then OK again.



For the exam remember that in order to communicate with systems off the network, your computer will need an IP address, subnet mask, and default gateway configured. To communicate with systems on the network you only need an IP address and subnet mask configured.

Configuring TCP/IP en masse using DHCP

If you're the network administrator of a large network, you don't want to run around to 400 workstations and configure an IP address, subnet mask, and default gateway on each computer. Not only is this time-consuming to initially set up, but it also becomes a nightmare to manage because of all the potential for human error. I have spent my days running around to each computer on the network, a sheet of paper in my hands, making sure that each computer is configured properly, and I can tell you that it is *not* fun!

Today's network operating systems support a feature called *Dynamic Host Configuration Protocol (DHCP)*. DHCP is a standard that allows the network administrator to tell the DHCP server a range of IP addresses that it is allowed to give out, along with the other TCP/IP options such as a subnet mask and default gateway. When the DHCP server has been configured to give out the addresses, the desktop computers automatically request an IP address from the server when they start up, and the server hands them all the IP address information. This means that the network administrator doesn't have to run around to each computer individually to configure TCP/IP, and in the long run, that saves time and money.

The steps to configure a Windows 2000 or XP system to obtain an IP address from a DHCP server are very similar to actually assigning the IP address manually.

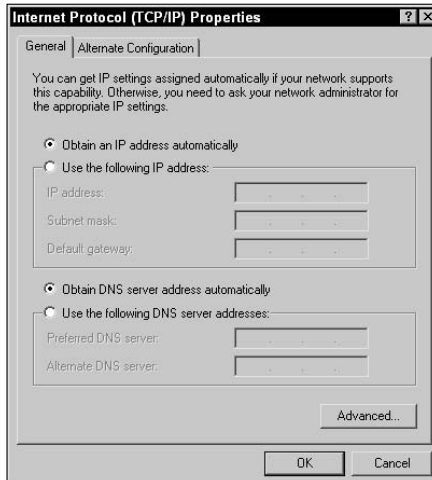
To configure a Windows 2000 or XP client for DHCP, follow these steps:

1. **If you're using Windows 2000, choose Start⇨Settings⇨Control Panel⇨Network and Dial Up Connections. If you're using Windows XP, choose Start⇨Control Panel⇨Network and Internet Connections⇨Network Connections.**
2. **Right-click your local area connection and choose Properties.**
3. **In the list of items used by the connection, select TCP/IP and click Properties.**

The Internet Protocol (TCP/IP) Properties dialog box appears.

4. **On the General tab, select the Obtain an IP Address Automatically option, as shown in Figure 3-9, to enable this system to be a DHCP client.**

Figure 3-9:
Configuring
Windows
XP as a
DHCP client.



5. Click OK and then OK again to close the network connections dialog box.

Automatic Private IP Addressing

If a DHCP server isn't available, and your Windows clients are configured to obtain an IP address automatically, will they receive an IP address? For operating systems before Windows 98, the answer would be no. But all versions of Windows since Windows 98 support a feature called *Automatic Private IP Addressing (APIPA)*. APIPA is a feature that allows the client to self-assign an IP address if the DHCP server doesn't respond to the DHCP request. The address that the client self-assigns is within the 169.254.x.y network range. The system will also configure itself with a subnet mask of 255.255.0.0 but will not configure the default gateway entry. This means that if the DHCP server is down and your network clients boot up, they will all have an address in the 169.254.x.y range and will be able to communicate with one another. Because they are not configured for a default gateway entry, they will not be able to communicate with systems off the network or with the Internet.



For the exam, remember that when you are troubleshooting networking connectivity, you use the `ipconfig` command to view the TCP/IP settings on a client. If the IP address is 169.254.x.y, then that means the client can not communicate with the DHCP server. Make sure the client is connected to the network correctly and then verify that the DHCP server is functioning!

Understanding Name Resolution

To communicate with another computer across a TCP/IP network, you have to know the IP address of the computer you are trying to communicate with. This is unrealistic considering that you are probably not too interested in trying to memorize all the IP addresses of the different Web sites you visit every day.

When running a TCP/IP network, you assign a friendly name to each computer and reference each computer by the friendly name instead of using the IP address. This means that instead of using an address like 204.56.78.6 to connect to Bob's computer, you would use a friendly name, like `bob`.

There are two types of names to understand when troubleshooting TCP/IP networks: *Computer names* (also known as *NetBIOS names*) and *fully qualified domain names*.

NetBIOS names

In the Windows world you access resources on a system by connecting to the computer name of the system. As an administrator, you assign a computer name to each computer on the network. The *computer name* (*NetBIOS name*) is a friendly name of up to 15 characters that is assigned to a computer and is used to uniquely identify the computer on the network. Users can then connect to the computer by the computer name *or* by the IP address — they will find it much easier to remember the computer name!

Changing the computer name in Windows 2000/XP/Server 2003

The process of changing your computer name in Windows 2000/XP/2003 is a common task to want to perform, and is fairly straightforward with today's Windows operating systems. To change your computer name, follow these steps:

- 1. Right-click My Computer and choose Properties from the context menu.**

If you don't see My Computer on the desktop, you can find it in the Start menu.

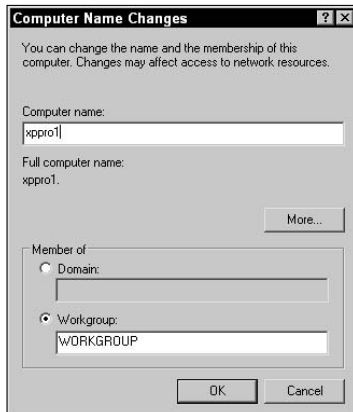
- 2. Click the Computer Name tab.**
- 3. Click the Change button.**

The Computer Name Changes dialog box appears, allowing you to type a new computer name, as shown in Figure 3-10.



When changing your computer name, notice that you can set the workgroup as well. A *workgroup* is the term given to a logical grouping of computers. When users browse the network, they may choose a workgroup, such as “Accounting,” and they will see any systems that are a part of the “Accounting” workgroup. To place your system in a particular workgroup, simply type the name of the workgroup in the text box.

Figure 3-10:
Changing
your
computer
name in
Windows XP.



4. Click OK to close out all the dialog boxes.

5. Reboot the system.

Because the computer name does not take effect until you reboot the system you are required to reboot.

WINS

When you network in a Microsoft environment, you connect to other computers by using those computers’ names. These computer names must be converted to IP addresses for communication to happen in a TCP/IP network. For example, you may want to connect to Bob’s computer, so you connect to `\\bob` through the Run command on the Start menu. When you try to connect, `\\bob` has to be converted to an IP address for the computer to be able to look for it. The process of converting a name from one format to another is *name resolution*. In this example, the computer name is being converted to an IP address, which is known as *NetBIOS name resolution*.

With NetBIOS name resolution, before your computer tries to connect to another system, it sends a query to a *Windows Internet Naming System* (*WINS*) server, asking the server this: “Hi there, Mr. WINS server. I am trying to connect to a computer named `bob` — do you have an IP address for this

computer?” The WINS server holds a database of NetBIOS names and matching IP addresses, known as the *WINS database*. Think of this database as having two columns: One for the computer name (NetBIOS name) and one for the matching IP address. Upon receiving the question, the WINS server checks the database for the computer named `bob` and then returns its IP address to the client who asked for it. Then the client can connect by using the IP address for `bob`.

In order for a Windows system to send a query to the WINS server you must ensure that you configure the WINS server setting within the TCP/IP properties of the client system. Configuring the Windows client for a WINS server directs the client to the server that it must register its name and IP address with and also whom to send name resolution queries to.

Before configuring your clients for WINS you must be aware of the IP address used by the WINS server. If you don't know this information you should consult the network administrator. Once you have the IP address of the WINS server, you are ready to configure the WINS clients. To configure a Windows 2000/XP/2003 system as a WINS client, follow these steps:

- 1. If you're using Windows 2000, choose Start⇒Settings⇒Control Panel⇒Network and Dial Up Connections. If you're using Windows XP/2003, choose Start⇒Control Panel⇒Network and Internet Connections⇒Network Connections.**
- 2. Right-click your local area connection and choose Properties.**
- 3. In the Network Properties dialog box, select TCP/IP and click Properties.**
- 4. In the Internet Protocol (TCP/IP) Properties, click the Advanced button.**
- 5. Click the WINS tab.**
- 6. Click the Add button to add the IP address of the WINS server.**
- 7. Type the IP address of the WINS server and click Add, as shown in Figure 3-11.**

The LMHOSTS file

If you don't have a WINS server, and an application you are running requires the use of a NetBIOS name (computer name), then you need to use what is known as the `LMHOSTS` file. The `LMHOSTS` file resides on each computer and is used to resolve, or convert, computer names to IP addresses. This file exists on each system on the network — you simply need to add an entry for the computer name and the corresponding IP address for each system you want the file to resolve. Figure 3-12 displays a typical `LMHOSTS` file.

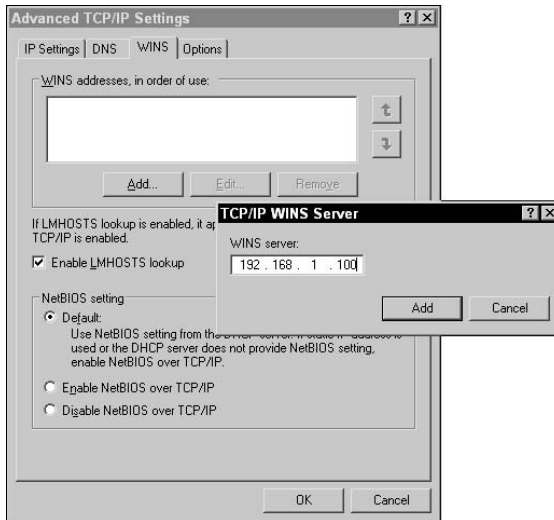


Figure 3-11:
Configuring
a WINS
client in
Windows XP.

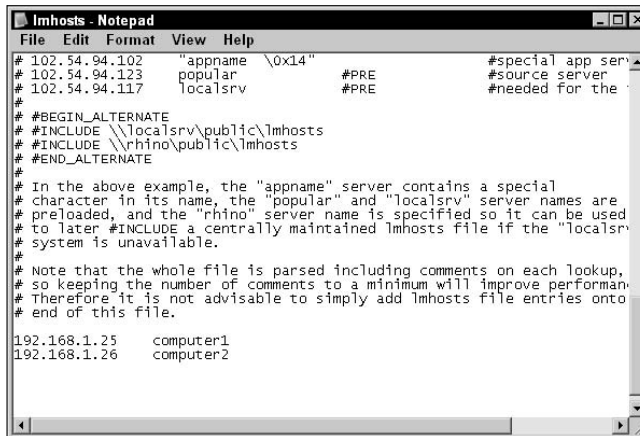


Figure 3-12:
An example
of an
LMHOSTS
file in
Windows XP.

Windows 2000/XP/2003 stores the file in %systemroot%\system32\drivers\ etc. In Windows 2000/XP/2003, the folder has an existing LMHOSTS file that you can use as a sample, but it has a .SAM extension that needs to be removed because the true LMHOSTS file has no extension.

Fully qualified domain names (FQDN)

The other type of name that can be assigned to the computer when you are running a TCP/IP network is a host name, or a *Fully Qualified Domain Name*

(*FQDN*). FQDNs are used when you run a TCP/IP- or Internet-based application like FTP, e-mail, or Web browser applications. For example, to navigate to my Web site via your favorite Web browser, you would type `www.gleneclarke.com` — this is an example of an FQDN. An FQDN is an Internet-style name that needs to be converted to an IP address in order for communication to occur.

The point is that when you use a computer name or an FQDN on a TCP/IP-based network, the names always need to be converted to the actual IP addresses. Again, the converting of names (either computer names or FQDNs) to IP addresses is the process referred to as *name resolution*.

There are a few techniques for FQDN resolution, and some are more popular than others. The following sections describe the name resolution techniques and their purposes.

DNS

DNS stands for *Domain Name System* and is the desired name resolution technique for resolving (converting) fully qualified domain names to IP addresses. Remember that FQDNs are the names that are used with Internet-based applications, such as e-mail and Web browsers. DNS is like a big database of FQDNs and their matching IP addresses. Think of this database as having two columns — one for the FQDN and the other for the IP address.

When you are running Internet or TCP/IP applications and you type in a fully qualified domain name, your computer sends a query, which is just a question, to the DNS database asking something like this: “I am trying to connect to `www.gleneclarke.com`. Do you have the IP address that matches this FQDN?” The database looks up the FQDN and returns the IP address to your computer, and your computer then connects to that IP address.

The big question is where is the database stored? The database is stored on what are called *DNS servers*. These servers are where the actual records are located and also where each client computer on your network sends its name queries.

To configure a Windows XP client to use a DNS server, you will add the IP address of the DNS server while configuring TCP/IP (refer to the section “Configuring TCP/IP in Windows 2000/XP/2003,” earlier in this chapter). Notice also that in the newer versions of Windows, the DNS server option is on the same screen as where you assign the IP address for a computer — it shows how critical DNS is to today’s computing! To configure a Windows client to use DNS, follow these steps:

1. In Windows 2000, choose Start⇒Settings⇒Control Panel⇒Network and Dial Up Connections. In Windows XP/2003, choose Start⇒Control Panel⇒Network and Internet Connections⇒Network Connections.
2. Right-click your Local Area Connection and choose Properties.
The local area connection Properties dialog box appears.
3. In the item list, select TCP/IP and click Properties.
4. In the TCP/IP Properties dialog box, select the Use the Following DNS Server Addresses option and type the IP address of your DNS server, as shown Figure 3-13.

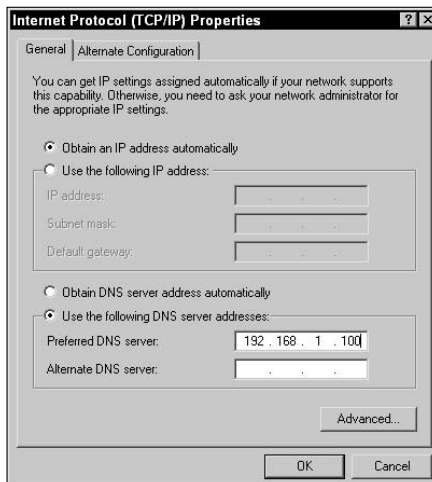


Figure 3-13:
Configuring
a Windows
XP client
for DNS.

5. Click OK twice to close the dialog boxes.

The HOSTS file

Very similar to the idea of using an LMHOSTS file for NetBIOS name resolution (see “The LMHOSTS file” section, earlier in the chapter), you can use a text file, called the HOSTS file, for host name resolution. The HOSTS file is located in the %systemroot%\system32\drivers\etc folder on Windows 2000, XP, and Windows Server 2003 systems.

The configuration of the HOSTS file is similar to the configuration of the LMHOSTS file; you simply create two columns — one for the FQDN and the other for the IP address of the system. You separate these columns with a Tab keystroke. Figure 3-14 shows a sample HOSTS file.

Figure 3-14:
An example
of a HOSTS
file used for
FQDN
resolution.

```

hosts - Notepad
File Edit Format View Help
# Copyright (c) 1993-1999 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#
# This file contains the mappings of IP addresses to host names. Each
# entry should be kept on an individual line. The IP address should
# be placed in the first column followed by the corresponding host name
# The IP address and the host name should be separated by at least one
# space.
#
# Additionally, comments (such as these) may be inserted on individual
# lines or following the machine name denoted by a '#' symbol.
#
# For example:
#
#       102.54.94.97       rhino.acme.com       # source server
#       38.25.63.10      x.acme.com           # x client host
#
127.0.0.1      localhost
192.168.1.235  server1.glensworld.loc
192.168.1.10  server2.glensworld.loc

```

ARP

Be aware that another layer of resolution needs to happen after your client has the IP address of the system it wants to communicate with. You know that there is a different database for FQDN resolution and computer name resolution — the DNS database stores fully qualified domain names and their IP addresses, while the WINS database stores computer names and associated IP addresses. Once the name is converted to the IP address, the IP address then must be resolved to the hardware address that is burned into the network card. This hardware address that is assigned to each network card is known as the *Media Access Control (MAC) address*. This means that there has to be a process that converts the IP address to the MAC address — a process known as *Address Resolution Protocol*, or ARP. ARP is an address resolution protocol that converts the IP address to the physical address assigned to the network card.

ARP is a *broadcast*, or a *yell*, out on the wire for a particular address. Looking at an example of Computer A trying to send information to Computer B, after Computer A has the IP address of Computer B (204.56.78.3), Computer A yells at the top of its lungs, “Hey, 204.56.78.3! What is your network card’s MAC address?” This yell runs along the network and eventually reaches Computer B, which responds with its MAC address. After Computer A has the MAC address of Computer B, it can then send the data to Computer B.



It is important to note that ARP messages are *broadcast messages*, and broadcast messages do not pass through routers. This doesn’t cause a problem because when you want to communicate with a system on a different network, your system sends the data to the *default gateway*, or *router*, and then the router sends the data off the network by ARPing the router on the destination network. In this example, your system communicates with the router, so it would ARP the router, not the destination system!

Real-world scenario: Planning TCP/IP settings

Determine what TCP/IP services are required to fit the scenario below. The answer follows the scenario.

Scenario: You are responsible for implementing a TCP/IP-based network across all of the office locations in your region. What TCP/IP options will you set, and which additional TCP/IP-based services will you implement?

Answer: You may want to set a DNS server to store the FQDN and IP addresses of the different servers on your network. You may also want

to run a WINS server for computer name resolution because it is a WAN environment.

After you set up the DNS and WINS servers, you will install a DHCP for each network segment so that you can hand out IP address, subnet mask, and default gateway information to the clients on the network. You will also hand out the IP address of the DNS server and the WINS server through DHCP so that the clients are fully configured.



For the exam remember that DNS and the `HOSTS` file resolve FQDNs to IP addresses, while WINS and the `LMHOSTS` file resolve computer names to IP addresses.

Troubleshooting with TCP/IP Utilities

After you have TCP/IP installed and configured and you have your TCP/IP network running, it is important to be able to troubleshoot the network. When problems arise on a Windows network, you can use some of the following commands to do your troubleshooting:

- ◆ IPCONFIG
- ◆ PING
- ◆ TRACERT
- ◆ NBTSTAT
- ◆ NETSTAT
- ◆ PATHPING
- ◆ NSLOOKUP

The following sections discuss these popular TCP/IP utilities that are used to troubleshoot TCP/IP connectivity.

IPCONFIG

If you are running Windows 2000 or Windows XP desktops, you can run the `ipconfig` (*IP Configuration*) utility, which shows you the current TCP/IP configuration of the Windows desktop, such as the IP address, the subnet mask, and the default gateway. If the computer is a DHCP client, `ipconfig` identifies the server that has given the IP address and also shows how long the IP address will be used by the client. Table 3-3 shows some of the switches supported by the utility `ipconfig.exe`.

Table 3-3 IPCONFIG Switches

Switch	Description
<code>/?</code>	Shows a list of switches supported by <code>ipconfig.exe</code> and a brief description of each switch.
<code>/all</code>	Shows all TCP/IP information — for example, DHCP lease period and the DNS server.
<code>/release</code>	Releases the current IP address information assigned by the DHCP server.
<code>/renew</code>	Requests new IP address information from the DHCP server.

For example, to use the `ipconfig` utility and view just the basic TCP/IP settings, you type the following at a command prompt:

```
ipconfig
```

But to view all the TCP/IP settings, such as your MAC address or the IP address of your DNS server, DHCP server, and WINS server, use the `ipconfig /all` command. You may also view when your lease time is up for the address you have been assigned when you use the `/all` switch on `ipconfig`.



Lab 3-1 will give you additional practice working with the `ipconfig` command. Lab 3-1 can be found in the `Labs.pdf` file in the `Author` directory of the CD-ROM.

PING

One of the most popular TCP/IP utilities is the `ping.exe` utility. *Ping* stands for *Packet Internet Groper* and is used to test whether your computer can communicate with a remote network device. If the ping test is successful, you get a ping response from the remote device; if it is not successful, the response will time out. The general syntax for using the ping utility is `ping <IP address>`, the IP address being the IP address of the network device you are testing.

When troubleshooting TCP/IP communication problems, it is important to understand the steps to find exactly where the problem occurs. Is the problem

in the computer you are using, in the computer you are trying to connect to, or in the default gateway?

Table 3-4 lists the order in which to ping each network device when trying to figure out at what stage the TCP/IP communication is failing.

Table 3-4 Troubleshooting Network Connectivity	
<i>Address to Ping</i>	<i>Description</i>
127.0.0.1	This IP address is referred to as the <i>loopback</i> address. It always tests your own workstation's TCP/IP software to ensure that it has loaded.
IP address of your computer	After you get a response from loopback, ping the IP address that has been physically assigned to your network card. To find out what IP address is assigned to your network card, use <code>ipconfig.exe</code> .
IP address of default gateway	If you get a response from your network card's IP address, ping the IP address of the default gateway. If you don't know the IP address of your default gateway, then run <code>ipconfig.exe</code> to view the address of the default gateway. Remember that the default gateway is the router, which is responsible for passing information on to other networks. If you are having trouble communicating with the router, you can't communicate with any devices off the network.
IP address of remote computer	After verifying that the default gateway is not the problem, if you know the IP address of the computer you wish to talk to (and chances are you don't because you are not the person who assigns the address to that computer), you can try pinging it. If you get a timeout at this step, you know that the computer you are trying to communicate with is the problem.

When you're troubleshooting a system that doesn't have Internet access, use the `ipconfig` utility to view the default gateway of the system; then ping the default gateway.



Lab 3-2 lets you practice changing your TCP/IP settings and work with the ping command. Lab 3-3 allows you to test your TCP/IP communication. Lab 3-2 and Lab 3-3 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

TRACERT

The `ping` utility is probably the most used TCP/IP utility, and rightfully so. The `ping` utility is a very useful utility for troubleshooting communication problems, but the `ping` utility tells you only whether your computer has communicated with the remote hosts; it does not tell you what path the

information took. This is where the *trace route* (`tracert.exe`) utility is useful. It is similar to the `ping` utility in the sense that responses are sent back to you if communication is established. The difference is that `tracert.exe` sends a response from every network it hits on the way, not just a response from the final destination. So `tracert.exe` shows you the path the information takes and also the number of networks between your computer and the computer you are talking to.

The trace route utility uses the following syntax:

```
tracert <IP address or DNS name>
```

NBTSTAT

Another popular network troubleshooting utility is `nbtstat`, which is used to troubleshoot NetBIOS name resolution. (Remember that discussion? If not, look back to the “NetBIOS names” section.) `nbtstat` stands for *NetBIOS over TCP/IP Statistics*. When your system resolves a computer name to an IP address, it stores that information in memory (known as the NetBIOS name cache) so that the next time the name needs to be converted to an IP address, the request is resolved from cache instead of broadcasting or querying a WINS server. If you wanted to verify that the entry is in cache, then you would use the `nbtstat` utility!

There are a number of uses for the `NBTSTAT` utility, so there are quite a few switches for the command. Some of the most useful switches are listed in Table 3-5.

Table 3-5		NBTSTAT Switches	
<i>Switch</i>		<i>Description</i>	
/?		Shows a list of switches supported by <code>nbtstat</code> and a brief description of each switch.	
-c		Displays the contents of the <code>NBTSTAT</code> cache. This cache shows the computer names and matching IP addresses that have been resolved recently.	
-A <ip address>		Displays the list of NetBIOS names used by the IP address typed with the <code>-A</code> . The listing also indicates what types of services the system is running.	
-n		Displays the NetBIOS names used by the local system.	
-r		Lists which addresses have been resolved through WINS.	

To use the `nbtstat` command, type something like the following at a command prompt:

```
nbtstat -A 192.168.1.200
```


Sharing File System Resources

In this section, I discuss how to set up the Windows operating system for the sharing of network resources. The topic begins with a discussion of some of the core options that must be set in order to network in a Windows operating system.

Every Windows computer must have a computer name in order to participate in a Windows network, and each computer name must be unique on the network in order to properly address a specific computer. To review the steps to change your computer's name, take a look at the section, "Changing the computer name in Windows 2000/XP/2003," earlier in the chapter.

Another important networking option to talk about before you take a look at sharing resources is the level of access control that is set on the computer. *Access control* determines what level of security you want to place on your Windows 9x desktop. Figure 3-15 shows the Access Control dialog box found in the network icon of the Control Panel on a Windows 9x system.

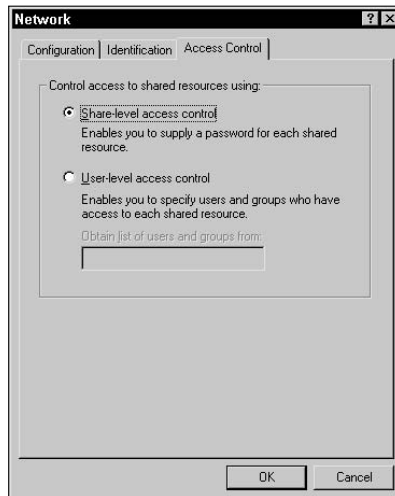


Figure 3-15: Configuring access control for shared resources in Windows 9x.

Share-level access control

There are two levels of access control to enable on a Windows 9x desktop, the first being share level. *Share-level access control* means that the security is placed on the share. A *share* is the concept of publishing a resource, like a folder or printer, out on the network so that others can connect to it. When you share a resource and share-level access is enabled, you have the opportunity to place a password on the share.

This is limited security because you are basing the security model on the fact that if someone knows the password, he or she can get into the share. It is limited because a user on the network could find out the password by accident and then be able to connect to the share without anyone knowing. Wouldn't it be better to assign a specific user the right to the share, so that only that particular user could get to the resource, thus increasing the security? That is just what the next section is about.

User-level access control

The second type of access control is user level. *User-level access control* means that when you share a resource out onto the network, you actually assign permissions to particular users. If someone wants access to the share, his or her user account has to have been given permissions to the share, and the user must log in with that particular account.



Windows 2000/XP/2003 products always use user-level access control. There is not an option to switch to share-level access control. This means that when you share a resource in these environments, you need to choose a list of users or groups who have access to the resource — this is a good thing!

Enabling File and Printer Sharing in Windows 2000/XP/2003

Now that you understand the two levels of access control, you are ready to allow your Windows machines to share resources on the network. First, you must ensure that File and Printer Sharing Services are installed and that File and Printer Sharing is enabled, and then you can start sharing folders and printers.

All NT-based products, such as Windows NT, Windows 2000, Windows XP, and Windows Server 2003, have File and Printer Sharing enabled by default. To verify that File and Printer Sharing is enabled within these operating systems, perform the following steps:

- 1. If you're using Windows 2000, choose Start⇨Settings⇨Control Panel⇨Network and Dial Up Connections. If you're using Windows XP/Server 2003, choose Start⇨Control Panel⇨Network and Internet Connections⇨Network Connections.**
- 2. Right-click your local area connection and choose Properties.**
- 3. Click the check box beside File and Printer Sharing for Microsoft Networks to enable it, as shown in Figure 3-16, and click OK.**

When this check box is selected, File and Printer Sharing is enabled, and you're finished. If File and Printer Sharing for Microsoft Networks wasn't listed, you need to install it first, so keep reading.

Figure 3-16:
Verifying
that File and
Printer
Sharing is
installed on
Windows XP.



4. If **File and Printer Sharing** isn't listed, click the **Install** button to install the service.
5. Choose **Service** in the **Component Type** dialog box and then choose **File and Printer Sharing for Microsoft Networks**. Then click **OK** to close all the dialog boxes.

Creating shared folders

When a user on the network wishes to access a file on another system, he must connect to a share on that system. *Shares* are a way to publish the folder on your system out to other users on the network so that they can access the files in that folder. If you have not shared any resources, then there is no reason for anyone to want to connect to your computer — it would be like giving someone the key to a locked but empty room.

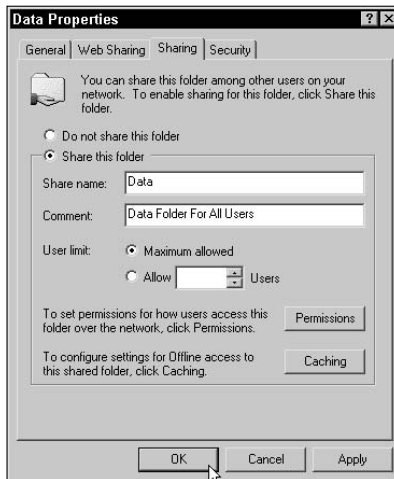


Know that you can only share folders or printers; you cannot actually share a file specifically. To allow users to access a file from across the network, you have to place the file in a folder and then share that folder.

Sharing a folder in Windows 2000

To share a folder on a Windows 2000 network, simply right-click the folder you wish to share and choose the **Sharing** command. In the **Sharing** dialog box that appears, select the **Share This Folder** option to share the folder. Within the **Sharing** dialog box, you will need to set a number of options (shown in Figure 3-17). The options are listed below:

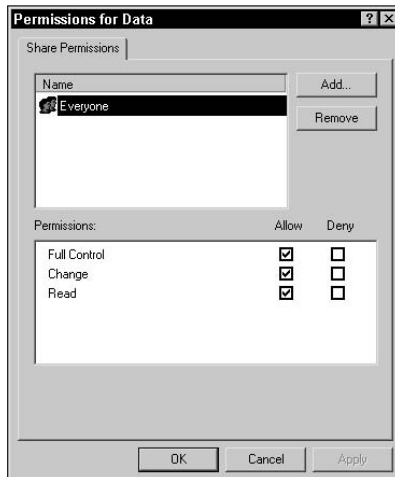
Figure 3-17:
Sharing a
folder in
Windows
2000.



- ◆ **Share name:** You need to give the share a name. This is the name that will be referred to by users who want to connect to the share.
- ◆ **Comment:** This is an optional description of the share that displays in Windows when the user views the list of shares in Detail view.
- ◆ **User limit:** You may limit how many users can connect to the share at any given time. This could be useful if you notice that the system is slow after a certain number of users connect. For example, if you're sharing a CD-ROM, you may notice that access to the CD-ROM slows after six users connect. In this example, you may want to set the user limit to five. Setting the user limit to maximum allowed will configure the user limit for 10 users connected to the share at once because the Windows desktop operating systems can only allow 10 connections at a time.
- ◆ **Permissions:** Clicking the Permissions button allows you to set permissions on the share. You set permissions to control which users can modify data in the share and which ones can simply read information in the share.
- ◆ **Caching:** This feature allows the client to store a local copy of data accessed in the share. This could be useful if you wish to allow a laptop user to take a copy of the data home and update the data. The modified data could then be synchronized with the content on the server when the user returns to the office.

Notice in Figure 3-17 that I have shared the folder named `Data`. As mentioned earlier, when sharing a folder, you need to set the share permissions. To set the share permissions, click the Permissions button. The default permissions in Windows 2000 are Everyone and Full Control, as shown in Figure 3-18.

Figure 3-18:
Setting
share
permissions
in Windows
2000.



This means that any user can create, delete, and modify files in the share or modify the share permissions. These default permissions are not good! The following steps show you how to adjust these permissions to your liking:

- 1. To remove the Everyone group from the permissions list, click the Remove button.**
- 2. Add specific users to the permissions list by clicking the Add button.**
The Select Users or Groups dialog box appears.
- 3. Select which user or group is allowed to connect to the share by selecting the user.**
You may add multiple users by clicking on the first user and then holding down the CTRL key and clicking on additional users.
- 4. Click Add to return to the Permissions dialog box.**
- 5. To set a user's permissions, select one of the following permissions for that user:**

There are three different permissions that you can assign to a user when dealing with shares:

- *Full Control:* The Full Control permission allows a user to read and change the contents of files on the share, to delete files on the share, and to change the share permissions. This permission is not normally assigned to users.
- *Change:* The Change permission allows a user to read the contents of files in the share, change the contents of files that exist in the share,

and delete files. Users cannot change permissions on the share with the Change permission.

- *Read*: The Read permission allows a user or group to read but not modify the content in the shared folder.

6. Click OK and then click OK again to exit the dialog boxes.

Sharing a folder in Windows XP

The steps to share a folder on a Windows XP system are very similar to sharing a folder in Windows 2000, but Microsoft has tried to simplify managing the security by hiding the security settings by default. To share a folder in Windows XP, follow these steps:

- 1. Locate the folder you wish to share on your system.**
- 2. Right-click the folder and choose Sharing and Security.**
- 3. On the Sharing tab, choose the Share This Folder on the Network option (shown in Figure 3-19).**

Notice that the share name is the same name as the folder.

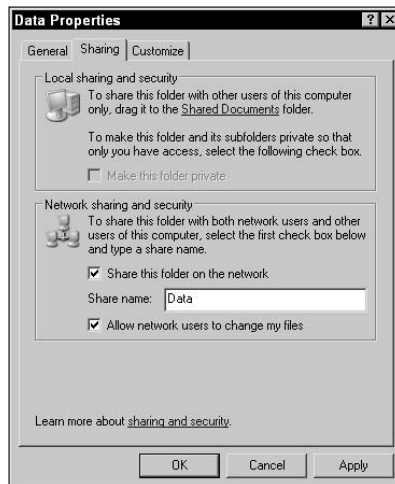


Figure 3-19:
Sharing a
folder in
Windows XP.

4. If you would like users to be able to alter the files in the folder, choose the Allow Network Users to Change My Files option.

5. Click OK.



You share a hard drive or CD-ROM in exactly the same way that you share a folder. To share the hard drive or CD-ROM, open My Computer, right-click the drive, and then choose the Sharing command — a great idea to share the CD to a number of users on the network so that all users can access the contents of the CDROM at one time!

If you want to modify the permissions on the share in Windows XP like you do with a shared folder in Windows 2000, you need to disable the *Simple File Sharing* option within Windows XP. To disable Simple File Sharing, follow these steps:

1. Open My Computer.
2. Choose Tools → Folder Options.
3. Click the View tab.
4. Scroll to the bottom of the Advanced Settings and clear the Use Simple File Sharing (Recommended) option, as shown in Figure 3-20.
5. Click OK.

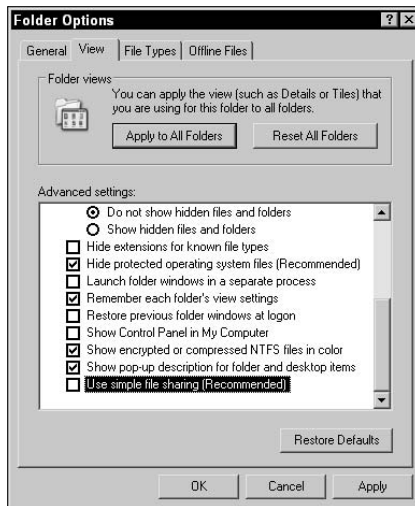


Figure 3-20:
Disabling
simple file
sharing in
Windows XP.

After you have disabled Simple File Sharing in Windows XP, you will be able to set the share permissions by selecting individual users and groups and what permissions you would like to assign by following the instructions in the “Sharing a folder in Windows 2000” section. Also note that Windows XP Home edition only supports Simple File Sharing.

Hidden shares

In the Windows world, you can also create hidden shares. *Hidden shares* are like normal shares in the sense that users on the network can connect to them; the difference is that hidden shares are not advertised — you can't find them by browsing through the shared folder list on a server. Users will connect to the hidden share by typing the UNC in the `Run` command or mapping a drive, which you can read about in the “Connecting to Shares” section.

To create a hidden share, use the steps for creating a normal share (see the section “Creating shared folders,” earlier in this chapter), but when you type a share name in the Share Name text box, you create the hidden share by appending a dollar sign (\$) to the end of the share name. For example, if the share name is `data`, and you want it to be a hidden share, you would type `data$` in the Share Name text box; the share is automatically hidden from Windows and users on the network when they browse the servers.

Multiple shares

In Windows 2000/XP/2003 you have the ability to create multiple shares for the same folder. This gives flexibility to the network administrator so that a user can have different permissions for a single folder, depending on what share that user connects to.



On our office network, we have implemented multiple shares per shared folder so that during day-to-day activities, not even an administrator can alter files on the server. If an administrator wants to make changes to a folder, he has to connect to the secondary share for that particular folder to have full-control access. This helps prevent a lot of unfortunate mistakes in modifying or deleting files by accident — even network administrators make mistakes! A big rule I follow is “protect the network from yourself as well!”

Connecting to shares

After you have created the shared resource, you can connect to the shared resource from anywhere on the network. There are a number of ways to connect to shared folders; here are a few of the most common:

- ◆ **Browsing My Network Places**
- ◆ **Using a UNC path through the `Run` command**
- ◆ **Mapping a drive**

The following sections examine each of these methods.

Browsing network resources

To browse network resources, follow these steps:

1. Go to My Network Places in Windows 2000/XP.
2. Click the View Workgroup Computers link on the left in the Network Task list.
You see a list of computers.
3. Double-click a computer to see a list of shares on that computer.
4. You can open any share just by double-clicking it.



Remember that you can't see any hidden shares while browsing network resources. For this reason, it is important to know additional ways to connect to shares, such as through the UNC path.

Using a UNC path

You may also connect to a share by using the *Universal Naming Convention (UNC) path*. The UNC path is made up of two backslashes (\\), the computer name you want to connect to, one more backslash, and the share name of the folder you want to connect to. The entire syntax looks like this:

```
\\computername\sharename
```

You would type this into the Run command, found by clicking the Start button.



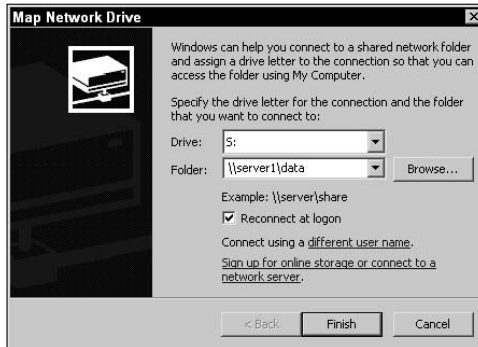
Using UNC paths means that you have to be aware of the exact names used for resources on the network, including hidden shares. When you get used to the computer names and share names on the network, you'll find that the Run command is quicker than waiting to see the list of computers in Network Neighborhood or My Network Places.

Mapping a network drive

You may also connect to shares by mapping drives. If you find that you are constantly connecting to the same resource, you may want to map a drive for the sake of simplicity. The idea of mapping a drive is that, in the end, you have a new drive letter in your My Computer folder that points to the UNC path of the resource. After the drive is mapped, anytime you wish to access the folder on the network, you go to My Computer and double-click the mapped drive.

To map a drive, right-click My Computer and choose Map Network Drive. In the Map Network Drive dialog box, select the letter for the drive you want to create and then type the UNC path to the shared resource into the Path text box. You may also choose the option to re-create this drive mapping the next time you log on so that you do not have to do this again. Figure 3-21 shows the Map Network Drive dialog box.

Figure 3-21:
Mapping a
network
drive in
Windows XP.



Lab 3-4 allows you to practice sharing resources and connecting to resources. Lab 3-4 can be found in the `Labs.pdf` file in the `Author` directory of the CD-ROM.

Sharing Printer Resources

You share printers in much the same way that you share folders on your system. After you have installed the printer and configured the settings so that the printer is functioning properly, it is time to share it.

Sharing a printer in Windows 2000/XP/2003

To share a printer, follow these steps:

- 1. Choose Start→Printers and Faxes.**
- 2. Right-click the printer you want to share and choose Sharing.**
- 3. Select Share This Printer and type the name of the share.**
- 4. Click OK.**

When sharing printers, all the same rules for sharing folders apply as far as the share name goes and how to create hidden shares.

Installing a network printer in Windows 2000/XP/2003

To print, or connect, to a shared printer out on the network, you have to install a network printer on your Windows client that points to the UNC path of the shared printer. A *network printer* in Windows is a printer installed that refers to a shared printer on the network. When you print to a network

printer the print job is sent to the computer that has the printer installed and prints from the print device connected to that system.

You can install a network printer in a number of ways. The two most popular methods are through the Add Printer Wizard and through the `Run` command, depending on the network setup. To install a network printer using the Add Printer Wizard you will run the Add Printer Wizard from the Printers folder. The wizard starts up and walks you through connecting to a shared printer. The following section shows you how to install a printer that points to a network location.

Installing a network printer is very similar in all Microsoft operating systems today. To install a network printer in Windows XP, follow these steps:

1. Choose **Start** → **Printer and Faxes**.
2. Click the **Add a Printer** link.
3. Read the welcome message and click **Next**.
4. Select the **network printer** option and click **Next**.
5. Choose **Connect to This Printer** and type the UNC path to the printer.
6. Click **Next**.
7. Click **Finish**.

Installing a network printer by using Point and Print

One of my favorite ways to install a network printer on a client is by taking advantage of Point and Print within Windows. Point and Print is a feature that copies the printer driver from one system to another as soon as you connect to the printer — no matter how you connect to the printer! What method is quickest to connect to the printer? Using the UNC path in the `Run` command!

After you type the UNC path of the shared printer you wish to connect to, Windows asks whether you want to install the printer on your system. When you choose **Yes**, a new printer is created in the Printers folder without your needing to run through the wizard. The printer driver is automatically copied from the system sharing the printer to your local system.

To install a printer by using Point and Print, follow these steps:

1. Choose **Start** → **Run**.
2. Type `\\computername\PrinterShareName`, where *computername* is the name of the system that is sharing the printer, and *PrinterShare Name* is the name of the shared printer.

3. When asked if you wish to install the printer, click Yes.

The printer is now installed and ready to print to.

Understanding Windows Services

In this section, I introduce you to the concept of a service within the operating system and then give an overview of some of the most popular services found in Windows.

A *service* is a software component within the operating system that provides a specific feature of the operating system. For example, the spooler service is responsible for providing printing functionality to the operating system. In order for you to print, you must have the spooler service running. When you're troubleshooting a system that isn't working properly, make sure that you check that the service that manages that aspect of the operating system is running. You may also want to restart the service if it is running, which is essentially a "rebooting" of that service.

Restarting a Windows service

To stop, start, or restart a service in Windows XP, follow these steps:

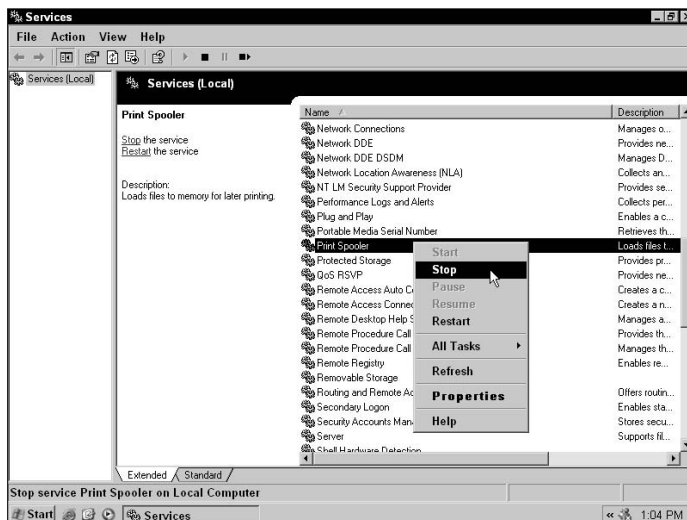
- 1. Choose Start⇨Control Panel.**
- 2. Choose Performance and Maintenance⇨Administrative Tools.**
- 3. In the Administrative Tools window, double-click the Services icon.**
- 4. After the Services console opens, you can stop, start, or restart a service by right-clicking the service and then choosing the appropriate command from the context menu, as shown in Figure 3-22.**

After you know how to stop and start a service, the next step is to understand some of the key services that exist within the Windows operating system and what the service offers to the system.

Server service

One of the critical services responsible for the networking of the Windows operating system is the Server service. The *Server service* provides file and printer sharing capabilities to the operating system. So, if you want to print to a printer that has been shared on Computer A from your system, Computer A needs to have the Server service running to allow your system to connect to it.

Figure 3-22:
Stopping a
service in
Windows XP.



Workstation service

The Workstation service is the exact opposite of the Server service. The *Workstation service* is responsible for making the connection to the system that is running the Server service. Going back to the example in the previous section, if you want to print to a printer on Computer A, the Workstation service on your computer sends the request to the Server service on Computer A.

DHCP Server service

The *DHCP Server service* handles DHCP server functionality, which is a server that hands out IP addresses to clients on the network. This service must be running if you want the DHCP server to hand out the IP addresses to clients on the network. This service exists only on Windows NT and 2000 Servers, and Windows Server 2003.

Print Spooler service

As mentioned earlier, the *Print Spooler service* is responsible for the printing environment in Windows. If this service is not running, you will be unable to print.

Messenger service

The *Messenger service* is responsible for sending and receiving messages within the operating system. For example, some features of the operating system allow an alert to be sent to the network administrator — the

Messenger service is responsible for the message (alert) being sent. You can send a message to another user on the network at any time by typing the following command into a command prompt:

```
Net send BOB "Hi there"
```

The above command sends a message to a user or computer called BOB, and the message that appears on Bob's screen says `Hi there`. In order for Bob to receive the message, your computer and Bob's computer need to have the Messenger service running.

Browser service

When connecting to network resources, clients need to know who is out there providing the resources. They can find this out by contacting the computer browser for their workgroup. The *computer browser*, or *browse master*, is a computer that maintains a list of servers that have File and Printer Sharing services enabled.

The first computer in the workgroup to start up that has File and Printer Sharing enabled becomes the computer browser for that workgroup. All other computers that start up and have File and Printer Sharing enabled advertise themselves to the browse master so that the browse master can update the lists of computers that are sharing resources on the network.

When Windows clients browse the network, they contact the browse master and ask for a list of servers on the network. The browse master returns the lists to the client, and then the client connects to the appropriate system to see a list of shares provided by that system.

It is important to note that you can control which system will act as the browse master. The browse master is determined by an election process in which each system submits criteria used to determine which computer should be the browse master. Generally, the system with the newest operating system wins. Also note that server operating systems will win an election over desktop operating systems.

If you would like to control which system can be a browse master, you can do this by changing a setting in the Registry on Windows 2000, Windows XP, and Windows 2003. It is important to understand how to designate a browse master because certain machines on the network just won't be adequate browse masters. Being the browse master takes up resources on the system, so if you have a low-end computer, you might not want it to ever be a browse master.

To configure whether a computer will be a browse master, go to the Registry Editor by typing `regedit` from the Run command and set the following registry setting:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Browser\
Parameters
```

Modify the `MaintainServerList` setting to one of the following:

- ◆ **Auto:** When selected, the computer has the potential to become a browse master, and the decision is made through election. Any computer with *Yes* set takes precedence over this system's *Auto* setting and becomes the browse master.
- ◆ **Yes:** When selected, this means that this computer should be the browse master for the workgroup. This option will take precedence over any system with an auto setting.
- ◆ **No:** When selected, this computer will never become the browse master for the workgroup.



Be very careful anytime you venture in to modify the registry — making the wrong change can cause the system to become unstable or even prevent it from starting up.

Getting an A+

This chapter discusses the different networking components that allow a Windows operating system to function in a networking environment.

- ◆ Four major networking components are required in Windows networks:
 - A network client
 - A network adapter
 - A common protocol
 - A service
- ◆ TCP/IP is the most popular networking protocol used today. When installing TCP/IP, you need to configure the following:
 - An IP address
 - A subnet mask
 - A default gateway
- ◆ The NetBEUI protocol is a non-routable protocol used on small networks.

- ◆ The IPX/SPX protocol is primarily used on Novell networks.
- ◆ In Microsoft networking environments, to allow users to connect to your hard drive or printer, you must install the File and Printer Sharing service.
- ◆ You may use the `ipconfig` on Windows 2000, XP, and Server 2003 to view your TCP/IP configuration.
- ◆ The `ping` command is used to send test messages to a remote system to verify that communication can take place on that system.

Prep Test

- 1** What networking component allows you to connect to a Windows Server?
 - A File and Printer Sharing services
 - B Client Service for Netware
 - C TCP/IP
 - D Client for Microsoft Networks

- 2** There is a shared folder named `public` on a computer named `Server1`. What is the syntax to connect to the shared resource using the UNC path?
 - A `\\server1\public`
 - B `\\server1\data`
 - C `\\data\server1`
 - D `\server1\data`

- 3** Which of the following IP addresses has a default subnet mask of `255.255.255.0`?
 - A `10.45.65.78`
 - B `132.107.2.34`
 - C `48.123.45.67`
 - D `216.83.24.56`

- 4** You have installed a new Windows XP computer in a Novell environment, and you are having trouble connecting to the Novell server. You have verified that you have the Client Service for Netware installed. What other networking component should you make sure is installed?
 - A Client for Microsoft Networks
 - B File and Printer Sharing
 - C The TCP/IP protocol
 - D The IPX/SPX protocol

- 5** You would like other people in the office to be able to access the printer that is attached to your Windows XP computer on your small network. Which networking component must you install?
 - A Client for Microsoft Networks
 - B Client Service for Netware
 - C File and Printer Sharing for Microsoft Networks
 - D NetBEUI

- 6** Which of the following is an example of a Class B IP address?
- A 164.34.56.8
 - B 12.45.76.2
 - C 202.34.65.32
 - D 125.67.6.7
- 7** What two properties of TCP/IP must be configured to communicate with other hosts on a small, local network?
- A IP address
 - B DNS server
 - C Default gateway
 - D Subnet mask
- 8** A user tries to connect to a shared resource called `data` on `server1` by using the UNC path `\\server1\data` but is unsuccessful. The user knows the IP address of `server1` and can successfully ping it. What is the problem?
- A The user should type `\\data\server1`.
 - B The client computer is not configured to query the WINS server.
 - C The client computer is not configured to query the DNS server.
 - D The user needs to install TCP/IP.
- 9** Which utility allows you to see the path that information may take when communicating with a remote system?
- A `ipconfig.exe`
 - B `tracert.exe`
 - C `winipcfg.exe`
 - D `arp.exe`
- 10** Which of the following TCP/IP settings are required to connect to the Internet? (Choose all that apply.)
- A IP address
 - B WINS
 - C Subnet mask
 - D Default gateway
- 11** Which resolution technique is used to resolve the IP address to a MAC address?
- A DNS
 - B WINS
 - C ARP
 - D `ipconfig.exe`

- 12** Which utility allows you to see the IP address information on a Windows NT computer?
- A winipcfg.exe
 - B ipconfig.exe
 - C tracert.exe
 - D arp.exe
- 13** Which service is responsible for automatically assigning IP address information to each computer on the network?
- A DNS
 - B WINS
 - C winipcfg.exe
 - D DHCP
- 14** You have verified that you have a network adapter installed and an appropriate client, but you cannot connect to Computer B across NetBEUI. What could be the problem?
- A You have the wrong IP address assigned to the computer.
 - B Computer B is not running NetBEUI.
 - C You should check the frame type on your computer.
 - D You need to install File and Printer Sharing to connect to another computer.
- 15** When you type `ipconfig.exe`, you do not see the IP address of the DHCP that has given you the IP address. What should you do?
- A Type `ipconfig /renew` at the command prompt.
 - B Type `ipconfig /all` at the command prompt.
 - C Ping the IP address of the DHCP server.
 - D Type `ipconfig /release` at the command prompt.
- 16** You have installed a printer and want to share it out to the network. How can you do this?
- A Type `net print` at the command prompt.
 - B Install the Client for Microsoft Networks.
 - C Install the IPX/SPX protocol.
 - D Right-click the printer in the Printers folder and choose the Sharing command.
- 17** Which service is responsible for converting `www.wiley.com` to an IP address?
- A DHCP
 - B DNS
 - C WINS
 - D ARP

- 18** You want to share the data folder as a resource to be used only by network administrators. How should you do this?
- A Hide the share by naming it `data$` and then set up the permissions so that only administrators have access. Tell the administrators the UNC path.
 - B Don't share the folder; tell the administrators to go to the local computer to access the resource.
 - C Set up the share permissions so that only administrators have access to the share.
 - D Share the resource and don't set any permissions.
- 19** You need to install software on all five computers on your small network. What should you do?
- A Place the software in a CD-ROM on one computer and share the CD-ROM drive out to the network.
 - B You must go to each computer to install the software because CD-ROMs cannot be accessed from across the network.
 - C Use DHCP to automatically install the software.
 - D Type `net install` at a command prompt.
- 20** Which of the following services is responsible for keeping track of the computers on the network that are sharing resources?
- A DNS
 - B WINS
 - C Computer browser
 - D DHCP

Answers

- 1** D. In order to connect to a resource, you must have the appropriate client running. In this example, you are logging into a Windows server, so you must load the Client for Microsoft Networks. *See "Network client."*
- 2** B. To connect to a shared resource on the network you will use the UNC path. The proper syntax for a UNC path is `\\servername\sharename`. *Review "Using a UNC path."*
- 3** D. 216.83.24.56 is an example of a Class C address, whose default subnet mask is 255.255.255.0. The other addresses are Class A and Class B, which have different default subnet masks. *Check out "Subnet mask."*
- 4** D. After you have Client Service for Netware installed, you have to ensure that IPX/SPX is installed and configured for the same frame type as the server you are trying to communicate with. *Peruse "IPX/SPX."*
- 5** C. File and Printer Sharing must be installed to allow someone to connect to your resources. *Take a look at "Enabling File and Printer Sharing in Windows 2000/XP/2003."*
- 6** A. The number in Class B addresses' first octet ranges from 128 to 191. The number in Class A addresses' first octet ranges from 1 to 126, and the number in Class C IP addresses' first octet ranges from 192 to 223. *Peek at "IP address."*
- 7** A, D. Because the computer is functioning on a small local area network where there is not a router, all you need to configure the functionality of TCP/IP is the IP address and subnet mask. Because no information leaves the network, you don't have to configure a default gateway. *Look over "The TCP/IP Protocol."*
- 8** B. The computer name, `server1`, must be converted to an IP address. WINS is the service that maintains a database that holds computer names and matching IP addresses. *Study "WINS."*
- 9** B. `tracert.exe` is the utility used to see the number of networks between you and the remote host. `ipconfig.exe` and `winiipcfg.exe` display the TCP/IP configuration. *Refer to "TRACERT."*
- 10** A, C, D. To participate on the Internet, you need an IP address, a subnet mask, and a default gateway. The default gateway is the IP address of the router that sends information off the network. *Examine "Configuring TCP/IP in Windows 2000/XP/2003."*
- 11** C. Address Resolution Protocol (ARP) converts the IP address to the network card address (MAC address). DNS and WINS convert different types of names to an IP address. *See "ARP."*

Networking the Operating System

- 12** B. `ipconfig.exe` is the utility run on Windows NT products to view TCP/IP configuration. `winipcfg.exe` is the utility on Windows 95 and Windows 98. Review "IPCONFIG."
- 13** D. Dynamic Host Configuration Protocol (DHCP) is a service on a server that automatically assigns IP address information to each computer on the network, saving the network administrator from having to manually configure each computer. Check out "DHCP Server service."
- 14** B. If you have a client and a network adapter installed, and they are working correctly, then the reason you can't connect across NetBEUI is because the person on the other end is using a different protocol. Peruse "NetBEUI."
- 15** B. Without switches, `ipconfig.exe` displays only the IP address, subnet mask, and default gateway. Use the `ipconfig /all` switch to view all TCP/IP configurations, including the DHCP server that has assigned your computer an IP address. Take a look at "IPCONFIG."
- 16** D. After you have installed File and Printer Sharing, you need to right-click the printer and choose the Sharing command. Peek at "Sharing Printer Resources."
- 17** B. DNS is responsible for converting fully qualified domain names to IP addresses. WINS converts the computer name to an IP address, DHCP is responsible for automatic configuration of TCP/IP, and ARP is responsible for converting IP addresses to MAC addresses. Look over "DNS."
- 18** A. The best way to be sure that no one except network administrators can access a shared resource is to share the resource and set the proper permissions. When you share the resource, though, it may be best to hide it from Network Neighborhood so that no one tries to sneak into it. Study "Sharing File System Resources."
- 19** A. One of the most efficient ways of installing software these days is to share the CD-ROM on one computer and have the other computers connect to the CD-ROM to install. This is one of the purposes of having a networked environment. Refer to "Sharing a folder in Windows XP."
- 20** C. The computer browser service maintains a browse list, which is a list of computers that have File and Printer Sharing enabled. Examine "Browser service."

Chapter 4: Configuring Internet Access

Exam Objectives

- ✓ Getting connected to the Internet
- ✓ Working with TCP/IP technologies on the Internet
- ✓ Configuring Internet settings with Internet Explorer
- ✓ Working with firewall devices to protect your computer

In this chapter, you examine how to connect to the Internet and discover some of the basics of how it works. This chapter is chock-full of terminology that you need to be familiar with, as well as different methods that are available to access the Internet. Because the Internet now plays such an integral part in many people's lives, as a CompTIA A+ Certified Professional, you should be familiar with where it came from and how to connect to it.

Understanding the Internet

The Internet was originally created and implemented by *DARPA (Defense Advanced Research Projects Agency)* — often referred to as just *ARPA* — in response to a U.S. Department of Defense request. At the time, the U.S. Department of Defense was concerned about its centralized communications network. Most communications were relayed through a central computing system, and damage to that system could stop computer communication. To avoid this problem, the U.S. Department of Defense employed another government agency (ARPA) with the responsibility of devising a new system.

The first thing that ARPA had to do was create some communication protocols that would allow computers to talk to each other in a new and non-standard decentralized manner, and it needed a small network on which it could test and develop the new protocols. It eventually interconnected four hosts, and that formed the start of the ARPANET.

The ARPANET (and the now the Internet) was defined by a series of standards that are currently being put forth by the *Internet Architecture Board (IAB)*, which represents the governing body of the Internet. These Internet standards are defined by IAB but discussed in *RFC (Request for Comments)*



documents. The first RFC defined how the initial hosts on the ARPANET would send to and receive data from each other.

All of the RFCs are available from www.rfc.net — there are currently over 4,600 of them. Technologies in RFCs sometimes make it into *STD (Standards)* documents, of which there are currently fewer than 70.

During the 1970s, protocols were created and evolved, allowing for support of more and more services over the ARPANET, and the number of members continued to increase. Most of the members of the ARPANET were locations that did research for ARPA — mostly universities and research centers. Terminal connections via Telnet were created, as was the fashion of the @ sign as a separator between the username and mail server for e-mail addresses. Emoticons were soon to follow starting with joke markers -) in the late 1970s, and the first Smileys in the early 1980s.

During the 1980s, several major events occurred:

- ◆ TCP/IP became the standard communication protocol.
- ◆ The term *Internetwork*, or just *Internet*, became the name of the network.
- ◆ Many operating systems standardized on TCP/IP.
- ◆ The National Science Foundation took over management of the Internet.
- ◆ The IAB was established to manage accepted standards on the Internet.

Also during 1980s, other countries started to join their national networks to the U.S. Internet, making it a true world-wide network. Several new application protocols were created to allow communication, such as *NNTP (Network News Transfer Protocol)*, and *IRC (Internet Relay Chat)*. This was also the decade that saw the first Internet *worm* released and named for its creator — it was the Morris Worm.

TCP/IP is the protocol suite that is used by the Internet, but it is composed of many different protocols that function at different levels of the network model. There are network protocols, transport protocols, and application protocols. In fact, the number of protocols is limitless. Many key protocols have already been discussed (for example, in Book VIII, Chapter 1, and Book VII, Chapter 3), but this chapter covers some of the others, such as

- ◆ **SMTP (Simple Mail Transport Protocol)**
- ◆ **POP3 (Post Office Protocol version 3)**
- ◆ **HTTP (HyperText Transfer Protocol)**
- ◆ **FTP (File Transfer Protocol)**
- ◆ **IMAP (Internet Message Access Protocol)**



Do not forget that TCP/IP is the communication protocol used on the Internet, and that it is actually a suite of protocols that cover all aspects of the communication process.

In the 1990s, the biggest single change to the Internet was caused by *CERN* (*Conseil Européen pour la Recherche Nucléaire*, or *European Council for Nuclear Research*) developing a new method of linking documents stored on different servers together. They modified SGML (Standard Generalized Markup Language), created a new language called HTML (HyperText Markup Language), and named the technology the *World Wide Web*. This single technology changed the face of the Internet, which went from a method of linking documents for information to being used for shopping, personal expression, media production and delivery, and untold other things.

Since the year 2000, technologies created for and used over the Internet have steadily risen to allow for remote access, collaboration, file sharing, streaming media, and *VoIP* (*Voice over Internet Protocol*). Government services and major corporations have all embraced the Internet as a major method of communicating with clients, and people in general have changed the way they communicate with each other through instant messaging. Internet gaming has become popular, with software companies getting steady revenue streams from monthly gaming fees. Today, if the Internet disappeared, many people would be at a loss to figure out how to communicate with one another.

Many countries now offer public access terminals for the Internet in convenient locations, such as libraries. Users are able to connect to the Internet through these public access terminals, Wi-Fi hot spots, and any number of other methods. Some cities have gone so far as to set up city-wide public wireless access.

The Internet evolved from something small and took small steps to become what we now know as the Internet. Even today, the Internet is constantly evolving, with new protocols created daily and new uses for the technology limited only by imagination. I wonder what the Internet will look like by the end of the next decade.

Using an ISP

Essentially, the Internet is a large, routed network, with technology similar to the networks that would be found in a large corporation. The main difference between the Internet and a corporate network, however, is that all the small networks that make up the Internet are joined together, while corporate networks are kept private. To gain access to the Internet, you must gain access to one of the networks that are connected to the Internet. These networks are

run by Internet service providers, or ISPs, which are in place simply to allow you to access the Internet. Figure 4-1 illustrates how a connection to the Internet works. In order to get connected to the Internet, the following would occur:

1. You connect your computer to your ISP's network by using a communication device like a modem or a router.
2. The ISP connects its network to its provider's network by using a router and a communication link like a leased Telco line.
3. Eventually there is a connection made to part of the Internet backbone, which allows connections to every network that is connected to the Internet.

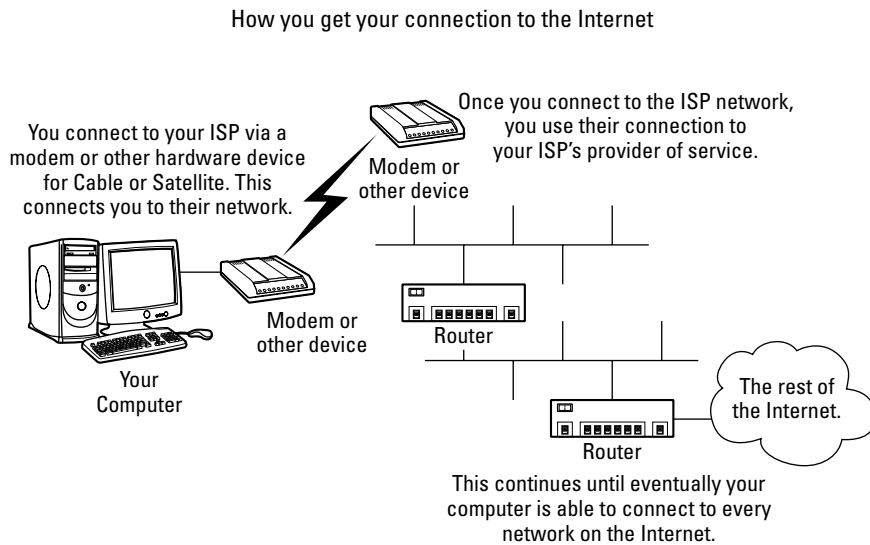


Figure 4-1:
How you get
connected
to the
Internet.

In addition to connection services, most ISPs provide other services, such as

- ◆ E-mail addresses and message space on servers.
- ◆ Local news servers that replicate USENET news groups. These are discussion groups for various topics in a bulletin board format.
- ◆ Web page storage space.

In the past, most Internet access service was provided over standard modems or through a dedicated link, such as 56K frame relay. Currently, dial-up access to the Internet is giving way to permanent connections like cable

and *ADSL (Asynchronous Digital Subscriber Line)*, which offer great improvements in access speed.

Working with an ASP

The acronym *ASP* is used for two different technologies: *Active Server Pages* and *Application Service Providers*.



You will not be tested on ASP for the exam. This information has been included because application service providers are sometimes considered to be a type of ISP, although they usually offer only data and content services and not connection services.



Active Server Pages are a form of HTML document that contains Visual Basic script or JavaScript that is rendered into straight HTML by the server before passing data to the Web client. They are used to limit the amount of work that has to be done by the client and to protect your code by sending only the results to the client.

Application service providers host servers for their clients. These servers may be database or mail servers. They may be shared or dedicated to a single customer. In either case, the server's hardware and software are maintained by someone other than the customer. This is a useful setup for smaller organizations because they do not have to invest in large servers or personnel to manage them — they basically rent space on them.

Understanding Internet Protocols

When many people think of protocols and computers, they think specifically of network and transport protocols — these two types of protocols seem to get all of the credit. Network and transport protocols are the network communications components that connect your computer to a network and send and receive bundles of data between hosts. They include protocols such as TCP/IP, IPX/SPX, and NetBEUI.

Generically, a *protocol* is a set of standards or conventions that are followed when formatting data to be used for electronic communications, and data transfer is just one level in the electronic communications model. This definition of protocol is not limited to data transfer, and there are a number of protocols that work at other layers, most notably the application layer. *Application-layer protocols* establish a standard or format for data that is to be communicated. They are called *application-layer protocols* because they are the first layer to which programs or applications on a computer (as well as the server components) communicate.

TCP/IP

TCP/IP (Transmission Control Protocol/Internet Protocol) is not a protocol in and of itself; rather, it is a suite of industry-standard protocols. It is a routable wide area network protocol that shares many similarities with Novell's IPX/SPX, which are both covered in Book VIII, Chapter 3. One main difference, however, between IPX/SPX and TCP/IP is that TCP/IP is an open, or free, protocol, while IPX/SPX is owned specifically by one company. TCP/IP standards are developed, established, and used by the computing community itself, while Novell is responsible for the development and standards for the IPX/SPX protocol.

Figure 4-2 shows some of the protocols that are used as part of the TCP/IP protocol suite and what each protocol is responsible for. At the lowest level, IP offers *best-effort delivery services*. This means that it attempts to deliver all network packets to the best of its ability. It also processes any errors that are reported back from routers. At the next level, TCP offers guaranteed delivery services, and UDP (User Datagram Protocol) offers best-effort delivery services.

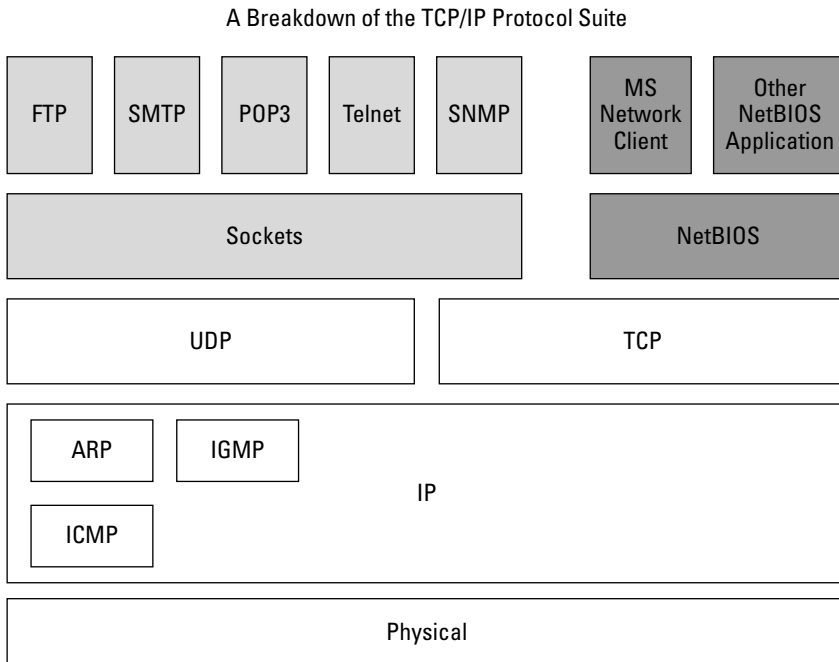


Figure 4-2: TCP/IP uses many protocols at different layers to accomplish its task.

Session services for TCP/IP are offered by either NetBIOS over TCP/IP (NetBT) or Windows Sockets (Winsock). The NetBIOS session interface is used by all Microsoft network clients including Windows 9x and Windows NT 4.0 as their method of communicating with Microsoft servers on the network. Windows 2000 and later Microsoft operating systems use the NetBIOS interface for backward compatibility, preferring to make connections using the Winsock interface. The Winsock interface is Microsoft's implementation of BSD Sockets, which is the primary session interface that has been used on all UNIX and UNIX-based systems. Since most Internet servers originally ran UNIX, the application layer protocols such as HTTP, FTP, SMTP, and POP3 are all designed to communicate through the socket-based session interface.

E-mail

E-mail is one of the applications that made the Internet indispensable to most people. E-mail was an early tool for ARPANET that enabled users to communicate ideas and concepts to colleagues many miles away. E-mail allows for individually addressed text messages to be transferred over the Internet and delivered directly to the targeted recipient(s). Compared to conventional land-based mail, these transfers are instantaneous.

Land-based mail is often referred to as *snail-mail* because of its slow speed compared to e-mail.

Attachments that accompany e-mail messages are converted into a text stream by means of encoding. *MIME (Multipurpose Internet Mail Extensions)* is currently one of the most popular encoding methods on the Internet. Other popular encoding formats include BinHex and Uuencode. *Encoding*, which takes 8-bits-per-byte binary data and converts it to 7-bits-per-byte ASCII or text data, enables binary attachments to be sent over the text-based e-mail network. When you receive attachments, they must be decoded by your e-mail program. If your e-mail program cannot do this, however, you have to use a third-party application to decode the files.

When it comes to reading your e-mail, you can choose from a wide variety of applications. In fact, the list of clients is extensive, including command-line clients, Windows-based clients, and Web-based clients.

Three basic protocols are used with e-mail: POP3, SMTP, and IMAP. Some e-mail clients may support more than one access protocol. SMTP and POP3 are the most commonly used protocols, although IMAP is increasing in popularity on private networks. Figure 4-3 shows how these three protocols fit together, and the following sections discuss them in more detail.

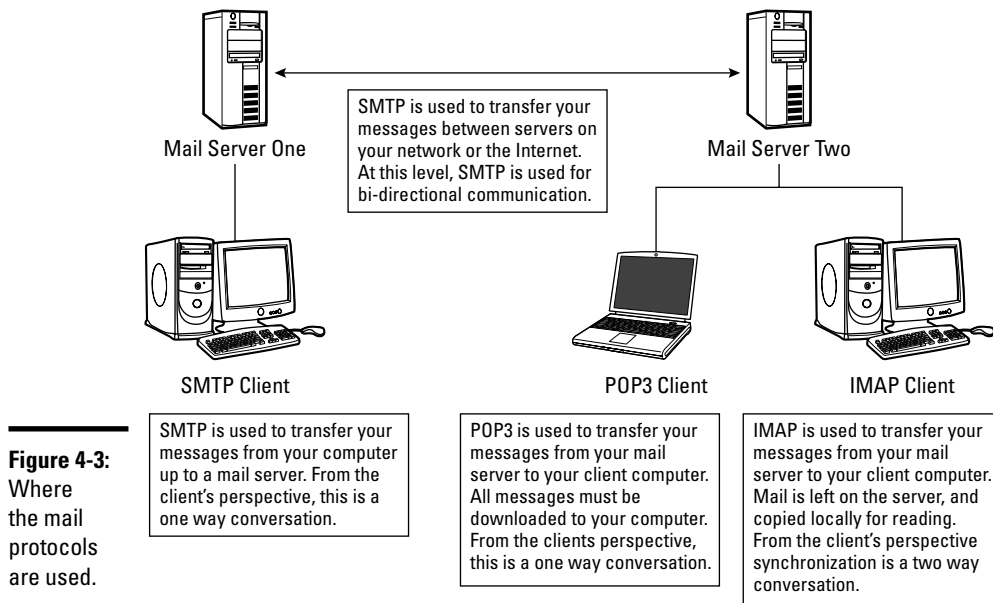


Figure 4-3: Where the mail protocols are used.

SMTP

The *Simple Mail Transport Protocol (SMTP)* is a mail delivery protocol. It is used to transfer mail messages from your mail client to a mail server. After the mail message is in the queue on the server, SMTP is also used to transfer the message to the mail server that is responsible for the target domain, such as `@mailtarget.loc`. The primary goal of SMTP is to get the mail messages to the targeted server.

POP3

The *Post Office Protocol version 3 (POP3)* is a client access protocol. POP3 is used to access or retrieve mail from a server. POP3 doesn't *send* e-mail — that's the responsibility of the SMTP. When you configure your e-mail client, you will configure it with the pair of servers, POP3 for downloading or reading, and SMTP for sending. POP3 clients usually download all mail messages for their servers and delete the mail from the server. This action then leaves the mail only on the client computer.

IMAP

The *Internet Message Access Protocol (IMAP)* is also a client access protocol for mail. As an IMAP client, you retrieve a list of messages that exist on the mail server and download only messages that you want to read. Any messages that

are downloaded are also left on the server. Changes to your mail files locally can also be replicated to the server. Because you can download all messages, you can work entirely offline. If you delete messages while working offline, those deletions will be replicated to the server the next time you connect, in turn deleting the messages on the server.



Know what the main e-mail protocols are and their role in the e-mail system.



HTTP mail accounts like hotmail and gmail are also becoming increasingly popular. These accounts leave their mail messages on the server in a manner similar to IMAP.

HyperText Transport Protocol

The concept of hypertext is much older than the Internet and has been around since 1945 when Vannevar Bush wrote an article named *As We May Think*. The word *hypertext* was coined in 1965 by Ted Nelson. Hypertext was a means of indexing or cross-referencing data found in different documents, allowing users to quickly move to linked documents. One commercial implementation of hypertext was the Apple Macintosh program Hypercard released in 1987, which could be thought of as a stack of virtual index cards, each holding common GUI items, such as text boxes, graphics, or buttons. These cards could be reviewed in order, or links could be made to allow for hypertext navigation through the stack. This feature allowed the Hypercard to be used for presentations, games, and a variety of other functions.

In the early days of the Internet, you could transfer data as files between computers. After being uploaded to the server by the data owner, these files were available for download from those servers, which meant you had to know what files you needed and what servers they came from. Your text files, formatted documents, and graphics could be transferred between computers. When you had downloaded the files, they could be opened and viewed, but there was no way to view them in an attractively formatted style in the online environment, especially in a format that was universally accessible. Most formatted documents were created and formatted in programs like Microsoft Word or Aldus PageMaker, which are readable only by people who have those programs.

In order to address the problem of knowing which servers you were accessing data from, Paul Lindner and Mark McCahill from the University of Minnesota came up with an idea that became the *Gopher protocol* in 1991. Gopher used hypertext concepts and allowed you to place a pointer on your server that would connect people to specific directories on other servers elsewhere on the Internet. This made browsing information that was scattered across servers very easy because switching between servers became completely transparent. Gopher became the most popular tool for downloading files and data from the Internet.

At the same time Gopher was being developed, the researcher Tim Berners-Lee at CERN (the European Organization for Nuclear Research) was working on a hypertext system he called World Wide Web (choosing that name over The Information Mesh). This system implemented a protocol called *HTTP*, which allowed for transparent linking of documents between servers. This capability was possible thanks to the new *HyperText Markup Language (HTML)* that was being used for World Wide Web data.

The tool that accesses HTTP servers and HTML files is, of course, the *Web browser*. The job of the Web browser is to retrieve the files from the server, display the formatted document, and link to other servers as required. Early Web browsers displayed text and images as separate but linked documents, so you could view images if you followed an image link. The Web really took off when, in 1993, Marc Andreessen of the National Center for Supercomputing Applications (NCSA) released NCSA Mosaic and offered something the Gopher and other Web browsers lacked — the capability to view text and graphics mixed together in a single frame. Internet users flocked to this new technology that gave their data the same appearance as the paper-bound copies. Figure 4-4 shows a formatted document in a Web browser.

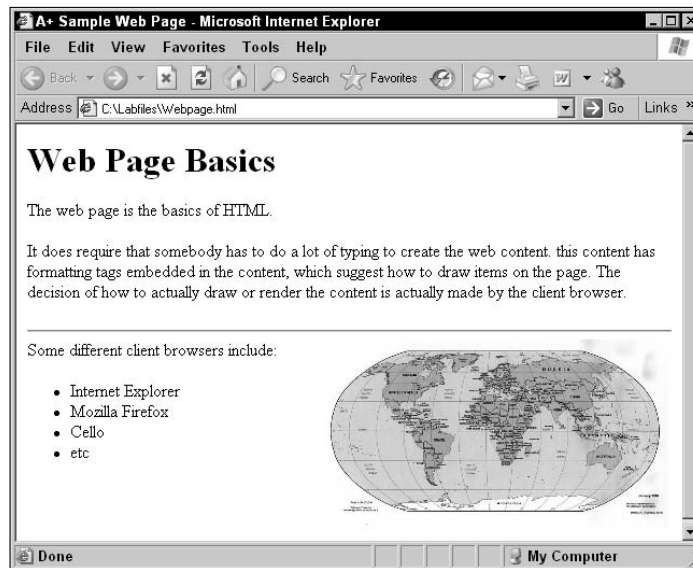


Figure 4-4:
Web browsers let you retrieve data from a server using HTTP.

HyperText Markup Language

HyperText Markup Language (HTML) is a form of *Standard Generalized Markup Language (SGML)* that offers a universal way to format documents. The standards for SGML are more complex than for HTML. SGML is actually a method of creating interchangeable structured documents, so that they

can be universally accessed from different types of systems. SGML can take data from a variety of sources like word processors and graphics applications, and join them together as a single structured document by using *Document Type Definitions (DTDs)*. Different DTDs are identified in a document with the aid of markups or tags which show or identify the divisions or sections. HTML is actually just a DTD that is one of the many small parts of SGML. With the simple initial requirements for formatting of documents, Tim Berners-Lee decided that keeping the formatting language simple was the best course of action for the Web, so HTML was adopted in favor of SGML.

HTML files are text or ASCII files, but they contain formatting codes that are embedded in the text. The Web, or HTML, page that is displayed in Figure 4-4 was generated with the following script:

```
<html>
<head><title>A+ Sample Web Page</title></head>
<body>
<h1>Web Page Basics</h1>
<p>The web page is the basics of HTML.</p>
<p>It does require that somebody has to do a lot of typing to
  create the web content. this content has formatting tags
  embedded in the content, which suggest how to draw items
  on the page. The decision of how to actually draw or render
  the content is actually made by the client browser.</p>
<hr>

Some different client browsers include:
<ul>
<li>Internet Explorer
<li>Mozilla Firefox
<li>Cello
<li>etc
</ul>
</body>
</html>
```

Early HTML files were saved with an extension of either `.html` or `.htm` (for MS-DOS 8.3 character compatibility), but now may have a variety of extensions, such as `.php` or `.asp`, due to various scripting languages that are being used to build dynamic Web pages.

HTML has gone through many revisions (the last revision was version 4). Changes are being made to HTML with the integration of *XML (Extensible Markup Language)* into HTML. XML adds functionality to support database data exchange. XML is designed to transfer not only sections of data from a database but also its structure. XML has become very popular and is the new Internet buzzword.

XML has been rolled into the HTML standard, forming the new *XHTML* (*Extensible HyperText Markup Language*) standard. So HTML is now XHTML, and the current version of XHTML is 1.0, with 2.0 just around the corner with a working draft already in place. For more information about XML, HTML, and XHTML standards, visit the World Wide Web Consortium at www.w3.org.

HTTPS and SSL

While HTTP is the HyperText Transport Protocol, *HTTPS* is a secure version of the protocol. When you're using HTTP, all communication between you and the server is in clear text, so if anyone wanted to, he or she could easily read all of the communication. HTTPS uses *SSL* (*Secure Sockets Layer*) or its successor, *TLS* (*Transport Layer Security*), to provide authentication and encryption services. While HTTP communicates with a server on TCP port 80, HTTPS uses TCP port 443.

In order to implement SSL, a Web server needs to have a security certificate installed to verify its identity for you. When you are confident that the server is who it says it is, you can then carry out secured and encrypted data transfers with it.

Early SSL used 40-bit encryption keys, which were not secure and were easy to break. Current SSL uses 128-bit or larger keys and is very hard to break, making data protected with SSL very secure.



Do not confuse the flurry of H-based acronyms used for the World Wide Web. HTML is the document formatting, HTTP is the clear text transfer protocol that uses TCP port 80, and HTTPS is an encrypted transfer protocol, secured by SSL or TLS, and which uses TCP port 443.

File Transfer Protocol

The first scientists using the Internet established early on during the evolution of the Internet that there would have to be some format to allow for the transfer of data (as opposed to text messages) across the Internet. The solution was *File Transfer Protocol* (*FTP*). FTP allows files to be uploaded and downloaded from servers. FTP requires a directory or folder that is going to be made accessible, a server-side service (or *daemon*), and a client. There are a large number of FTP clients to choose from, ranging from command-line to graphical. Figure 4-5 shows some different FTP clients that are available on the Windows platform.



A *Windows service* is an application that runs without a user interface and without anyone being logged onto the computer. This background application provides functions to the local computer or to remote users. These functions include items such as file- and printer-sharing, messaging, and publishing Web pages. Linux and UNIX servers call these background services *daemons*.

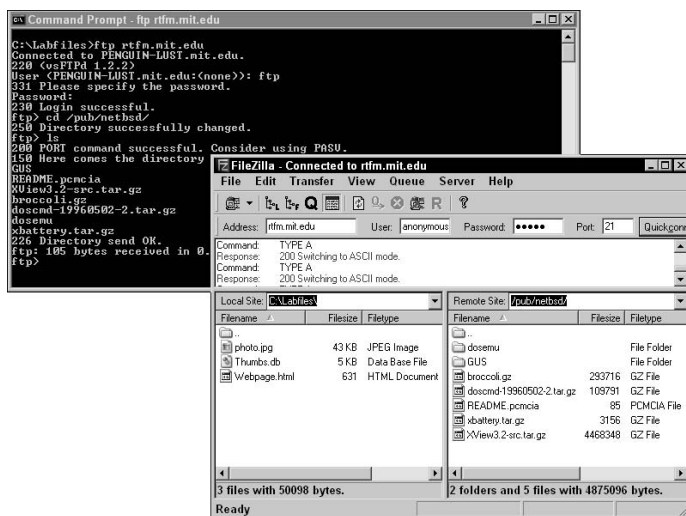


Figure 4-5:
Choosing
the right
client for
the job is
important.

Command shells

A *shell* is an interface for users to connect to. When you log onto your Windows computer, your Desktop environment is generated by `explorer.exe`. On UNIX and Linux servers, if administrators choose not to load a graphical environment, the shell application would be command-line driven, usually with some variant of *bash* (*Bourne Again Shell*). Because most UNIX or Linux server administration can be done from the command line, remote shell access to a server is very important, both from the aspect of remote administration and for security. In addition to these servers, many network devices, such as routers, switches, and hardware firewalls, support a remote management command shell. The two most common methods of connecting to these command shells are `telnet` and `ssh`.

telnet

`telnet`, developed in 1969, has been the long-standing standard for remote management, in spite of its long-standing security flaw. The largest single flaw with `telnet` is the fact that the entire communication process between the client and server takes place in clear text. For any person in a position to view the raw network traffic, the entire conversation can be viewed, including the logon usernames and passwords. Server administrators moved to newer management technology some time ago, but many of the hardware manufacturers kept using `telnet` until fairly recently. Windows includes a default command-line `telnet` application, and you can also get third-party graphical `telnet` applications that let you easily manage multiple sessions or capture log files of sessions.

ssh

`ssh`, developed in 1995, is short for *secure shell*. Remote access to a shell on a computer is what `ssh` is designed to secure, just like the name suggests. In addition to allowing remote shell access to servers, `ssh` allows for opening secure tunnels, allowing secure access to other systems on a remote network through the tunnel. When working with `ssh` with the lowest security level, the server uses a certificate to generate initial encryption keys and then encrypt the data stream between the client and server. For stronger security, client certificates can be used for authentication of the client as well as the server using its certificate. Server administrators that dealt with the `telnet`'s clear text data for years quickly converted to `ssh` when it became available, making it the new standard in remote shell access. Some `telnet` holdouts, like network switch manufacturers, have finally listened to their customers and added support for `ssh`.



`ssh` should be considered the minimum connection mechanism for remote devices, and `telnet`, with its clear text data transfers, should no longer be considered.

Installing and Configuring Browsers

Internet Explorer is the default Web browser for most Windows computers, partly due to the fact that it comes installed with operating systems. With this boost behind them, Microsoft won the browser wars of the 1990s, which saw the main battle between Internet Explorer and Netscape Navigator (both owed their roots to NCSA Mosaic). Although many different browsers were available in the beginning, these two products quickly rose to the top of the pile and battled it out.

Right around the turn of the century, Netscape released the source code for the core Netscape engine, calling it *Mozilla*, to the open source community. This powerful engine quickly formed the core of many open source products, and one product gained widespread adoption by the open source market — Mozilla Firefox. Firefox offers many features that were not originally included in Microsoft's products, such as pop-up blocking and tabbed browsing, plus the ability to choose from thousands of add-ons, called extensions, written by the open source community. This product was released at a time when Microsoft's Internet Explorer was experiencing several security problems, and many people wanted a product that did not suffer from the same problems.

As with all browser wars (and all wars), products are continually enhanced, and Microsoft's next version of Internet Explorer will include many of the features that have lured people to Firefox. You can find Firefox at www.mozilla.com/firefox, and the latest version of Internet Explorer is at www.microsoft.com/windows/ie. Both of these sites have links you can click to download and install the browser, and you will have the option each time you start the non-default browser to set it as your system default browser.

The following sections look at Microsoft's Internet Explorer version 6.0, although similar settings are available in most Web browsers.



To configure your Internet Explorer settings, open IE and choose Tools⇨ Internet Options, which opens the Internet Options dialog box shown in Figure 4-6. This dialog box contains seven tabs, and I discuss each of them in the following sections.

Figure 4-6: All of your Web browser settings can be managed through Internet Options.



Configuring General settings

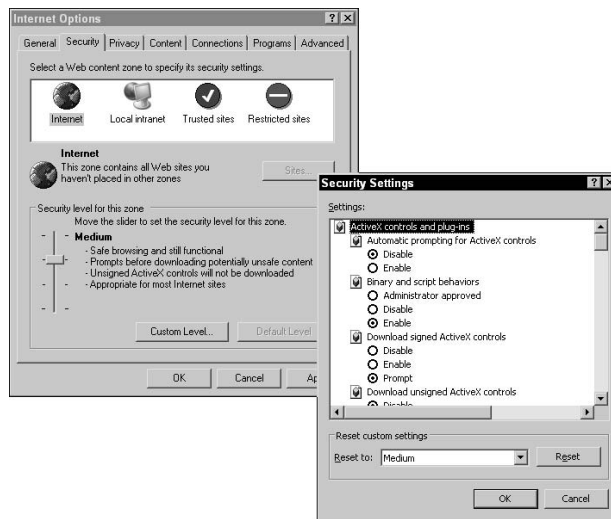
The General tab of Internet Explorer's Internet Options dialog box, shown in Figure 4-6, lets you configure settings for your Home Page, Temporary Internet Files, and your browser's History. The other settings are primarily cosmetic. The settings for Temporary Internet Files allow you to choose the location of the files, how much space they will take on your hard drive, and your setting for refreshing content from the Web site directly.

Configuring Security settings

The Security tab of Internet Explorer's Internet Options dialog box lets you set security settings for sites in one of four zones. Most sites that you visit will fall into the Internet Zone, while the Local Intranet, Trusted Sites, and Restricted Sites will affect only sites that you specifically enter onto those lists. Each of these zones can have different security settings. For each zone, you can choose from one of four security levels or specify your own custom settings. The four built-in levels are High, Medium, Medium-Low, and Low security.

If you click the Custom Level button, you can see exactly what settings are enabled. The settings in the dialog box, shown in Figure 4-7, let you tell Internet Explorer how to handle a variety of components that may be embedded in Web pages. You can have Internet Explorer disable, enable, or prompt you when it encounters certain components on Web sites, such as .NET Framework–reliant components, ActiveX controls and plug-ins, downloads (for files and fonts), scripting (which includes ActiveX scripting and Java applets), user authentication (which is often done automatically in the background), and a slew of user interface options.

Figure 4-7:
Security
Settings in
Internet
Explorer can
make the
browser
much safer.



Configuring Privacy settings

The Privacy tab deals with cookies and pop-ups and works in a similar fashion to the Security tab. It allows you configure settings for cookies on a per-site basis as well as configuring general settings by using either the slider for levels from Allow All Cookies to Block All Cookies, with four levels in between, or by clicking the Advanced button and customizing your cookie settings.

Cookies are settings that are stored by your Web browser, and are sent to a Web server when you return to a Web site. Web sites use cookies to track users, especially if they have a shopping cart transaction system. Shopping carts allow you to browse a catalog, select items you want, and check out. This mechanism is not only used at shopping sites, but also at download sites, where you are able to select multiple files, and download them all in one action at the end of your visit. In addition to shopping carts, many sites will store a unique ID in a cookie on your computer, so that they will be able to track your return to their site. In some cases, this is used to set user preference, or just to track unique visitors. Cookies are generally safe; no one

has yet found a way to implant malware in a cookie, and they take up only a small amount of memory. When visiting one Web site, you may see content that comes from another Web server, which is the case with most of the banner ads that appear on Web sites. If this other Web server attempts to have a cookie stored, it is called a third-party cookie. You might want to block third-party cookies or all cookies if you are concerned about having companies track your activity on their Web sites.

If you choose to block pop-ups, by clicking the Settings button, you can configure settings as well as configuring a list of sites that are allowed to use pop-ups.

Configuring Content settings

The Content tab allows you to enable the Content Advisor which makes use of voluntary tags that Web content developers can include on Web pages. These tags rate the level of language, nudity, sex and violence that is found on the page. With the Content Advisor enabled, you can block content that is above the customized configured level. The Certificates section allows you view and modify SSL certificates that Internet Explorer uses. The Personal Information section allows you to modify the AutoComplete settings that are used on forms you encounter on Web sites.

Configuring Connections settings

A proxy server acts as a middleman when you are requesting information from the Internet, allowing network administrators to restrict access out through the network firewall, as well as audit what Web sites network users are going to. If you need to use a proxy server to browse the Internet, you can use the *proxy* settings to configure your access to proxy servers. Proxy settings are found on the Connections tab of IE's Internet Options dialog box and can be configured for each of the dial-up settings individually or for your LAN connection. Both have similar settings, so I discuss the LAN settings, which you reach by clicking the LAN Settings button. The settings are shown in Figure 4-8 and include automatic proxy detection and configuration, which requires specific configuration settings on your network. In addition to this setting, you can manually configure a proxy server address and port for your particular network and use that same server for *http*, *https*, *ftp*, *gopher*, and *socks*. The *socks* setting allows you to proxy protocols other than the default ones that are listed. You might use *socks* to proxy Internet Relay Chat (IRC) or POP3. As with most cases, you can also configure a list of servers that are exceptions to using the proxy server, which may be the case when you have servers on your internal network that you don't want to use a proxy server for.

Configuring Programs settings

The Programs tab of IE's Internet Options dialog box lets you specify which programs you would like to use for different types of Internet services, such

as HTML editing, e-mail, and newsgroups. In addition, this tab also allows you to configure and manage Internet Explorer add-ons, which are third-party components that are loaded automatically by Internet Explorer.

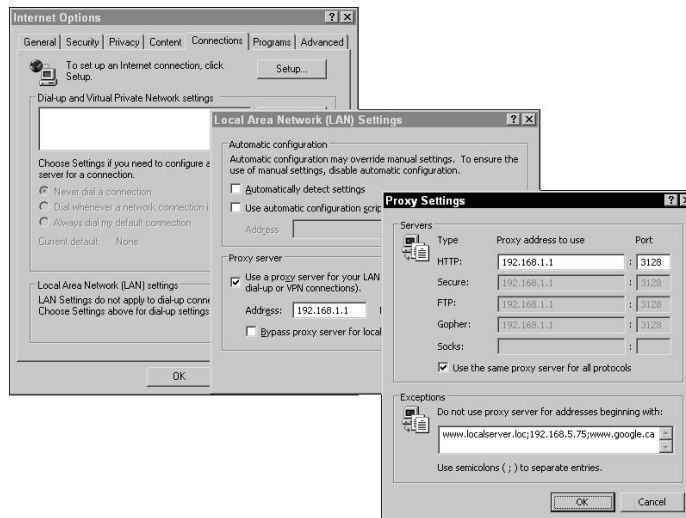


Figure 4-8: If you are using proxy settings, you may need to regularly enable or disable the settings.

Working with Advanced settings

The last tab in Internet Explorer's Internet Options dialog box is the Advanced tab. It has a variety of settings that let you modify, by enabling or disabling features and functions, how Internet Explorer works. These settings fall into the following categories: Accessibility, Browsing, HTTP 1.1, Java, Multimedia, Printing, and Security. Some of these settings are shown in Figure 4-9.

Using the Internet

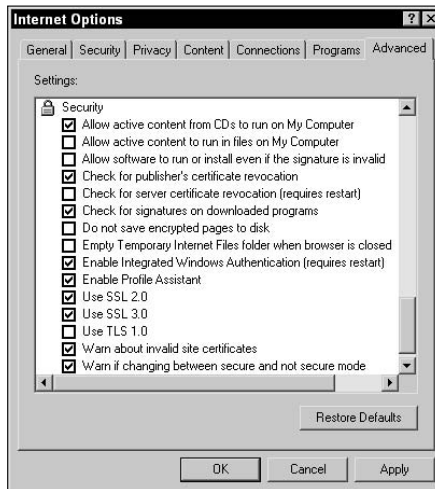
When you're using the Internet, you should be aware where you're going to when you click links on your Web browser. This section looks at what makes up the parts of a *URL (Uniform Resource Locator)*, the string of text that appears on the address line in Internet Explorer. The standard URL has an access method, a server name, and the path to a file or directory on the server. A typical URL resembles the following:

```
http://www.edtetz.net/sample-files/default.htm
```

In this example, `http:` represents the access method, `//www.edtetz.net` represents the name of the server that is being contacted, `/sample-files`

represents the directory that you're navigating to, and `/default.htm` represents the file that is being requested. All URLs follow the same basic structure, but if you leave out a part of it, like the document name, you will be given the default document for that directory or server.

Figure 4-9:
Security Settings are just one of the types of Advanced Settings.



Access methods

Table 4-1 summarizes some of the access methods for resources on the Internet. These access methods can be specified from the command line or from within a Web browser.

Table 4-1 Access Methods	
<i>Access Method</i>	<i>Description</i>
<code>http:</code>	<code>http:</code> is used to access Web content on servers.
<code>https:</code>	<code>https:</code> is secure <code>http:</code> access using <i>SSL (Secure Sockets Layer)</i> to provide authentication and encryption services. <code>https:</code> requires a certificate to provide verification of the server's identity.
<code>ftp:</code>	<code>ftp:</code> is used to copy files to and from remote servers.
<code>telnet:</code>	<code>telnet:</code> is used to access remote terminal services with a remote server.
<code>gopher:</code>	<code>gopher:</code> is a hyperlink protocol similar to <code>http:</code> but with less flexible display options.
<code>mailto:</code>	<code>mailto:</code> activates the local e-mail client to send a message to the address specified in the server portion of the URL.
<code>news:</code>	<code>news:</code> is used to access files on an <i>NNTP (Network News Transport Protocol)</i> server or USENET newsgroup.

Domain names and Web sites

Every computer on the Internet has a unique IP address. This address enables a computer to find and establish communications sessions with any other computer — as long as you know the IP address of the computer you want to connect to.

Because every IP address is a 12-digit number (and most people cannot remember hundreds of 12-digit numbers), the *Domain Name System (DNS)* was established in 1984. With this system, you can specify a name, and the DNS resolver or client on your computer will look up the required address for you on a DNS server.

Because each server would never be able to hold the names of all the computers on the Internet, DNS servers split up the job. Each server is responsible for knowing only a small number of computers, but the servers know how to find other servers. Figure 4-10 illustrates how the DNS is structured. At the top of the structure is the root (.) domain, which knows about all of the servers that manage the top-level domains. Top-level domains include `com` (commercial), `org` (organization), `mil` (military), `edu` (education), `gov` (government), and `net` (network), as well as a two-letter domain for every country in the world, such as `au` (Australia), `us` (United States), `uk` (United Kingdom), and `de` (Germany). The servers at the top level know about the servers that are responsible for the next level down, and so on.



DNS is a global directory that allows friendly names to be resolved to IP addresses. Without this function being performed, you would have to know every server's IP address.

If you had to connect to a server such as one with a DNS name of `www.edtetz.net`, your DNS client would check against your local server to see if it knows the IP address for that DNS name. Even if your server is not the owner of that DNS name, it may have looked it up before and have the information cached. If the information is not cached on the local server, your server will forward the request to one of the root-level servers, which will direct it to the `.net` server, which will direct it to the `edtetz` server, which will then look up the `www` record.

This system gives your browser the ability to find any computer on the Internet that has a name that was registered with DNS. Many people register their servers by the type of service that they offer, such as `ftp.edtetz.net`, `smtp.edtetz.net`, `pop3.edtetz.net`, or `mail.edtetz.net`. HTTP allows you to place links on any server to any other server on the Internet, which allows you to have a very complex path configured to lead people through the Internet. Because this path is web-like, the interconnection of HTTP servers is referred to as the *World Wide Web (www)*. `www` is the standard name that is given to the HTTP or Web servers on the Internet.

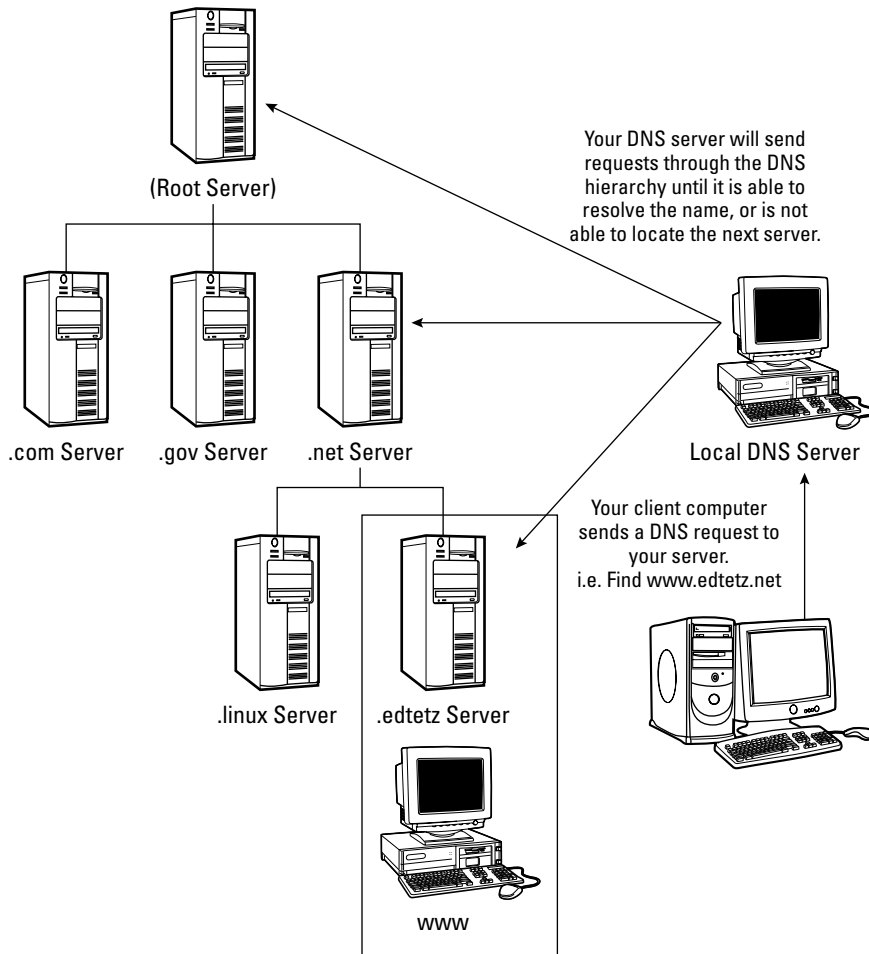


Figure 4-10:
Overview of the Domain Name System.

When you connect to a server, the URL may have a path listed beyond the root directory or may have the complete path to a file on the server. If no filename is given, the Web server usually displays a default document for that directory such as `default.htm` or `index.htm`. If there is no default document for that directory and the server allows for viewing the directory listing, then you will see a list of filenames that are found in the directory, with each filename being a link to the file. When no other condition is met, the server will return an Object not found or 404 error message.

Ways to Access the Internet

In addition to dial-up connections, many other types of connections can connect you to the Internet — some of these may be available in your area from local Internet service providers (ISPs). Some of these are used for home connections, and others are used primarily for connecting LANs to the Internet:

- ◆ **Dial-up**
- ◆ **Cable**
- ◆ **ADSL**
- ◆ **ISDN**
- ◆ **T1/T3**
- ◆ **Satellite**
- ◆ **Wireless**

The following sections take a closer look at each of these methods.

Dial-up

Good old reliable dial-up has remained a true friend of remote access and access to the Internet for years, mainly because for years it was the method that most users had available and dial-up access was much better than no access. It is still a standard way to connect to the Internet in many markets in which faster alternatives do not exist. The listed data rate for dial-up connections is 56 Kbps.

Cable

Cable connections implement a cable modem in your house that takes a digital network signal from your network card and translates (modulates) it into an analog broadband signal. This signal is then passed on to the cable network. Cable companies usually offer transmission speeds between 4 and 10 Mbps to their customers, while actual speeds may vary. This signal runs over the existing cable using previously unused signal areas of the medium. When using cable, you are on a shared medium with other users until your connection reaches the cable company's office.

ADSL

The *Asynchronous Digital Subscriber Line (ADSL)* implementation looks similar to that of cable, except the device you have in your home takes the digital signal from your network card and passes it on to a digital phone line. ADSL companies usually offer transmission speeds between 4 and 8 Mbps,

while actual speeds may vary. Although cable offers the same transmission speeds going to and from the Internet, ADSL always has slower upload speeds because the connection is broken into upstream and downstream channels. This means that surfing the Internet and copying large files from Web sites is very fast; but if you want to store a file on a Web site or send e-mails with large attachments your speeds will be substantially slower. Typical upload speeds for ADSL range between 64 Kbps and 1 Mbps.

The standards for ADSL2 and ADSL2+ have been released, so some vendors may choose to implement these versions of ADSL, allowing for download speeds from 5 to 12 Mbps for ADSL2 and up to 24 Mbps for ADSL2+, while both only offer upload speeds in the 1–3.5 Mbps range. With ADSL, you share the line only until you hit the telco switching office, which typically must be closer than 5 kilometers or 3 miles.

Broadband high-speed Internet connections through cable and DSL (of which ADSL is a subtype) have replaced dial-up in many regions.



ISDN

The *Integrated Services Digital Network (ISDN)* service comes in two basic forms: basic rate and primary rate. Basic rate ISDN uses three channels: two 64 Kbps lines for data (128 Kbps) and one 16 Kbps line as a control channel which is used for establishing and maintaining connections. The data channels are referred to as *B channels*, and the control channel is referred to as a *D channel*. Primary-rate ISDN uses twenty-three 64 Kbps B channels (1.44 Mbps) and one 64 Kbps D channel.

T1/T3

T1 connections offer transmission speeds of 1.544 Mbps over 24 pairs of wires. Each pair of wires can carry a 64 Kbps signal called a *channel*. T1 connections can be implemented over copper wire.

T3 connections, on the other hand, require a better medium than copper, such as microwave or fiber optic. They are capable of speeds ranging from 6 Mbps to 45 Mbps.

T1 and T3 are North American standards, while E1 and E3 are similar standards for the European community.

Satellite

Satellite Internet services come in two basic flavors: one-way with terrestrial return and two-way. Because satellites providing service in most residential areas were designed to send data to your house, the first Internet access

over satellites involved downloading from the satellite, but data had to be uploaded over a dial-up modem. This still tied up your phone lines, and your upload speeds were still rather slow. Two-way systems added technology to return a signal to the satellite. Speeds for uploads ran in the neighborhood of only 1 Mbps, but they freed up the phone lines. Download speeds over satellite systems rival those of broadband services like ADSL and cable. Just like satellite television, satellite Internet is susceptible to the weather and elements.

Wireless

Wireless access to the Internet is provided through standard 802.11 (Wi-Fi) wireless networks, which are set up to provide coverage in prescribed areas. In some locations, wireless access is provided for free to customers of certain businesses, or by a municipal government like Fredericton, New Brunswick (<http://www.fred-ezone.ca>). Sometimes a company may set up access points in a wide range of locations, and offer access through them as part of a subscription service. Wireless access points that are part of a subscription service will allow you to connect to them, but will only allow you to go to the company's Web site to set up an account. Once you have set up and paid for an account, you will be able to use any of the company's access points to access the Internet.

For more information about 802.11 networks, read Book VIII, Chapter 2.

Firewalls and Home Routers

When you are dealing with the Internet, it is important to remember that there are a lot of bad things out there, and you need some protection from the dangers of the world. All software firewalls have the inherent problem that they are software running on the system that they are trying to protect, which is considered bad because software (including firewalls) is more susceptible than hardware firewalls to breaches caused by Trojans or other viruses.

Home routers are typically *NAT (Network Address Translation)* gateways, which allow multiple computers to share one public IP address. This gives you a layer of additional security since it will only allow remote servers or computers to communicate with your computer, after you have initiated a connection to the remote computer. When you open the connection to the external server, a temporary channel opens to that server's IP address, allowing communication.

Some of these routers have more advanced firewall features to restrict the TCP and UDP ports that are allowed to be used between the two hosts, as well as the ability to filter or restrict access to URLs based on a series of

rules. Figure 4-11 shows some of the security settings that are available on these devices.

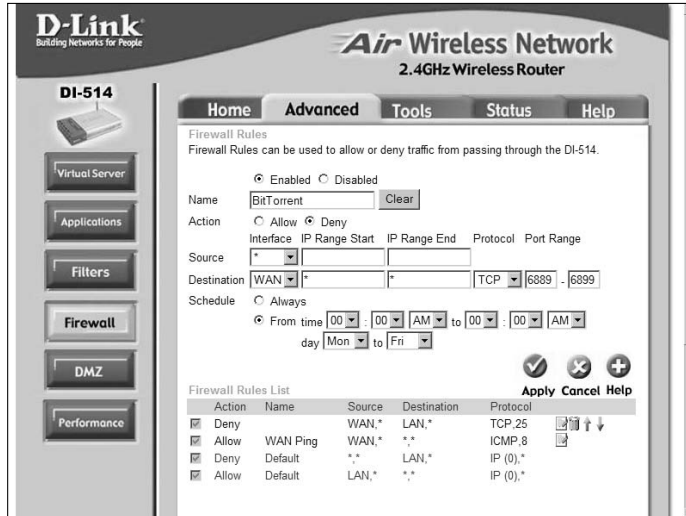


Figure 4-11:
Typical
firewall
settings on
a home
router.

In addition to these hardware residential and *SOHO* (*Small Office/Home Office*) routers and firewalls, there are various high-end hardware firewalls to choose from, as well as many Linux firewalls that boot directly from CD-ROM. For many of the Linux firewalls, you need a low-end computer with two network cards and a CD-ROM drive. For a home or small office, even a Pentium or Pentium II processor (or lower if you can find one) will provide enough power for the job.

Getting an A+

This chapter provides a brief history of the Internet and how you connect to it. Other concepts that are covered include

- ◆ TCP/IP is the suite of protocols that is used on Internet.
- ◆ The TCP/IP suite of protocols is made up of other protocols such as ftp, http, telnet, ssh, tcp, udp, and ip.
- ◆ POP3, SMTP, and IMAP are the primary protocols used for e-mail.
- ◆ HyperText Markup Language is the formatting system used for Internet Web pages.
- ◆ Dial-up, ISDN, T1, ADSL and cable are the main methods used to connect to the Internet.

Prep Test

- 1** What is the governing body of the Internet?
 - A W3.org
 - B DodNet
 - C IAB
 - D IABnet

- 2** What is the networking protocol used by the Internet?
 - A TCP/IP
 - B IPX/SPX
 - C Banyan/Vines IP
 - D HTTP

- 3** What organization provides Internet connectivity?
 - A W3.org
 - B ASP
 - C IAB
 - D ISP

- 4** What does the acronym ASP stand for? (Choose all that apply.)
 - A Active Server Pages
 - B Action System Performance
 - C Additives for System Providers
 - D Application Service Provider

- 5** The TCP/IP networking protocol is composed of how many sub-protocols?
 - A 2
 - B 5
 - C 10
 - D Many, too numerous to list

- 6** What purpose does MIME serve?
 - A It is used as an encryption method for e-mail.
 - B It encodes binary data into ASCII data so that it may be sent through e-mail.
 - C It is used to convert HTML data into an e-mail format.
 - D It is used to trap people inside invisible boxes.

- 7** What protocol is used to send e-mail?
- A POP3
 - B IMAP
 - C HTTP
 - D SMTP
- 8** What protocols are used to read e-mail? (Choose two.)
- A SMTP
 - B POP3
 - C IMAP
 - D UDP
- 9** Which of the following is used to format Web pages?
- A HTTP
 - B WWW
 - C SSL
 - D HTML
- 10** What protocol is used to transfer files to and from a remote server?
- A HTML
 - B POP
 - C FTP
 - D SMTP
- 11** What is the first part of an Internet URL (for example, `http://`) called?
- A Pointer
 - B Access method
 - C Control source
 - D Activation header
- 12** What system converts names you type in URLs into IP addresses?
- A Dual Naming Standards
 - B Name Resolution Server
 - C Name Recognition System
 - D Domain Name System

Configuring Internet Access

- 13** Which of the following are ways to connect to the Internet? (Choose all that apply.)
- A HSI/Net
 - B Cable
 - C Digital Phone Link
 - D ISDN
 - E Wire
 - F T1
 - G IAB
- 14** What does the acronym POP stand for?
- A Post Office Protocol
 - B Private Outside Postal Service
 - C Portable Office Protocol
 - D Peoples Office Protocol
- 15** HTML files are what type of files?
- A Text/ASCII
 - B Run-length encoded
 - C Binary
 - D Hyperlinked

Answers

- 1** C. The Internet Architecture Board (IAB) is the managing body of the Internet. See *"Understanding the Internet."*
- 2** A. TCP/IP is the Internet's networking protocol. Review *"TCP/IP."*
- 3** D. ISP stands for Internet Service Provider. Check out *"Using an ISP."*
- 4** A, D. The acronym ASP is usually short for Active Server Pages, but the term *Application Service Provider* is becoming very popular. Peruse *"Working with an ASP."*
- 5** D. Some of the sub-protocols that make up TCP/IP include SMTP, POP3, SNMP, FTP, TFTP, and 500 to 1,000 other protocols. Take a look at *"Understanding Internet Protocols."*
- 6** B. MIME (Multipurpose Internet Mail Extensions) is used to convert binary files into something that can traverse the text-only e-mail system. Peek at *"E-mail."*
- 7** D. SMTP is used to send mail messages over the Internet. Look over *"SMTP."*
- 8** B, C. POP3 and IMAP are used to read mail that is on a mail server. Study *"E-mail."*
- 9** D. The acronym HTML stands for *HyperText Markup Language*, which is the formatting language or standard for Web pages. Refer to *"HyperText Markup Language."*
- 10** C. The acronym FTP stands for File Transfer Protocol, which is the traditional method to transfer files to or from a server. Examine *"File Transfer Protocol."*
- 11** B. The first part of a URL represents the access method that is being used to connect to the server. See *"Using the Internet."*
- 12** D. The Domain Name System resolves requested names into IP addresses to allow you to connect to the named computer. Review *"Domain names and Web sites."*
- 13** B, D, F. Cable, ISDN, and T1 are methods of connecting to the Internet. Check out *"Ways to Access the Internet."*
- 14** A. The acronym POP stands for Post Office Protocol. Peruse *"POP3."*
- 15** A. HTML files are text or ASCII files. Take a look at *"HyperText Markup Language."*

Book IX

Securing Systems

The 5th Wave

By Rich Tennant



“A centralized security management system sounds fine, but then what would we do with all the dogs?”

Contents at a Glance

Chapter 1: Fundamentals of Security	1013
Chapter 2: Implementing Security	1039
Chapter 3: Viruses and Malicious Software	1065

Chapter 1: Fundamentals of Security

Exam Objectives

- ✓ Types of attacks
- ✓ Physical security
- ✓ Authentication and authorization
- ✓ Data protection

One of the most important skills to have if you're going to support networked systems or systems connected to the Internet is the fundamental skill involved in securing systems and networks. If you aren't working in a networked environment, you can apply these same skills to your customers with home Internet machines. The bottom line is that you need a solid understanding of network security in today's day and age.

I remember when a close friend of mine had his Web site totally replaced by a hacker. My friend's Web site files were replaced with inappropriate content, and he wondered how on earth someone had hacked his server. It seems amazing now, but back then (which was around 1994), a lot of companies didn't use firewalls because they weren't aware of the risks involved in having a computer connected directly to the Internet. Back then, people thought, "I have a password on the administrator account, so I am secure."

In this chapter, I introduce you to the basic concepts and terminology used to help secure an environment. Be sure to read this chapter carefully and make sure you understand the topics. Have fun with this topic area — it is very exciting!

Identifying Types of Attacks

If I had to define *hacker*, I would say that a hacker is someone who has the technical expertise to bypass the security of a network or operating system. A hacker also knows how to use features of a piece of software or hardware to gain access to restricted areas of a network and how to use those features against you and your system. For example, Windows 2000 servers have Web server software installed by default, and with that Web server

installed, anyone in the world can view or delete files on the hard drive of the server — and the hackers know this!

There are two types of hackers:


- ◆ **White-hat hackers** try to “hack” or break software or hardware for the purpose of understanding how to protect the environment from black-hat hackers. These are the good guys.
- ◆ **Black-hat hackers** are people who break into a system or network for malicious reasons or for personal gain. The reasons could be for financial gain, bragging rights, or revenge.

Hackers use a number of different types of attacks to hack into a network or operating system. Sometimes one attack can lay the groundwork for a future or different type of attack, meaning that the initial attack doesn't seem all that dangerous, but it is used in the future to gain unauthorized access. This section outlines some of the most popular types of attacks that can happen in networking environments today.

Social engineering attacks

A *social engineering attack* occurs when a hacker tries to obtain information or gain access to a system through social contact with a user. Typically, the hacker poses as someone else and tries to trick a user into divulging personal or corporate information that allows the hacker access to a system or network.

For example, a hacker could call your company's phone number, which is listed in the phone book, and pretend to be technical support for your company, telling the user who answers the phone that a new application has been deployed on the network and in order for the application to work, the user's password must be reset. After the password is reset to what the hacker wants, he may “verify” with the user the credential that the user uses. A user who isn't educated on social engineering may divulge important information without thinking.



A social engineering attack is an attack performed by a hacker where he tries to trick a user or administrator into divulging sensitive information through social contact. Once the sensitive information is obtained, the hacker can then use that information to compromise the system or network.

This example may sound unrealistic, but it happens all the time. If you work for a small company, you might not experience a social engineering attack, but for large companies, it is extremely possible that a social engineering attack would be successful if the company doesn't educate its users. A large company usually has the IT staff or management located at the head office, but most branch locations have never talked to IT management, so those

branch employees won't recognize the voices of the IT folks. A hacker could impersonate someone from the head office, and the user at the branch office would never know the difference.

There are a number of popular social engineering attacks scenarios — and network administrators are just as likely to be social engineering victims as “regular” employees, so they need to be aware. Here are some popular social engineering scenarios:

- ◆ **Hacker impersonates IT administrator:** In this example, the hacker calls or e-mails an employee and pretends to be the network administrator. The hacker tricks the employee into divulging a password or even resetting the password.
- ◆ **Hacker impersonates user:** In this example, the hacker calls or e-mails the network administrator and pretends to be a user who forgot her password, asking the administrator to reset her password for her.
- ◆ **Hacker e-mails program:** The hacker typically e-mails all the users on a network, telling them about a security bug in the operating system and that they need to run the `update.exe` file that is attached to the e-mail. In this example, the `update.exe` is the attack — it opens the computer up so that the hacker can access the computer.



Educate your users never to run a program that has been e-mailed to them. Most software vendors, such as Microsoft, state that they will never e-mail a program to a person — they will e-mail the URL to an update, but it is up to the person to go to the URL and download it. A great book to learn more on the process a hacker takes to compromise a systems is Kevin Beaver's *Hacking For Dummies, 2nd Edition*.

Network-based attacks

A *network-based attack* uses networking technologies or protocols to perform the attack. There are a number of different types of network-based attacks, the most popular of which are mentioned in the following sections.

Password attacks

There are a number of different types of password attacks. For example, a hacker could perform a *dictionary attack* against the most popular user accounts found on networks. With a dictionary attack, hackers use a program that typically uses two text files:

- ◆ One text file contains the most popular user accounts found on networks, such as administrator, admin, and root.
- ◆ The second text file contains a list of all the words in the dictionary, and then some.

The program then tries every user account in the user account file with every word in the dictionary file, attempting to determine the password for the user account.

To protect against a dictionary attack, be sure employees use strong passwords that mix letters and numbers. This way, their passwords aren't found in the dictionary. Also, passwords are normally case sensitive, so educate users on the importance of using both lowercase and uppercase characters. The hacker would not only have to guess the password, but also the combination of upper- and lowercase characters.

Also remind users that words found in *any* dictionary are unsafe for passwords. This means avoiding not only English words, but also French, German, Hebrew . . . even Klingon!

Hackers can also perform a *brute force attack*. With a brute force attack, instead of trying to use words from a dictionary, the hacker uses a program that tries to figure out your password by trying different combinations of characters. Figure 1-1 shows a popular password-cracking tool known as LC4. Tools like this are great for network administrators to audit how strong their users' passwords are.

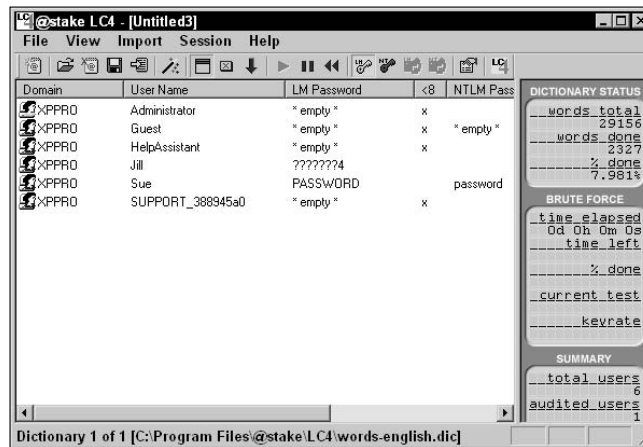


Figure 1-1:
Cracking
passwords
with LC4.



Remember that to protect against password attacks users should use strong passwords, which is a password made up of letters, numbers, symbols, has a mix of upper- and lowercase characters, and a minimum length of 6 characters.

Denial of service

Another popular network attack is a *denial of service (DoS)* attack. A denial of service attack can come in many forms and is designed to cause a system

to be so busy that it cannot service a real request from a client, essentially overloading the system and shutting it down.

For example, if I have an e-mail server, and a hacker attacks the e-mail server by flooding the server with e-mail messages, causing it to be so busy that it cannot send anymore e-mails for me, then I have been denied the service the system was created for.

There are a number of different types of DoS attacks. For example, there is *the ping of death*. The ping of death is when a hacker continuously pings your system, and your system is so busy sending replies that it cannot do its normal function.

Spoofing

Spoofing is a type of attack in which a hacker modifies the source address of a network packet. A *packet* is a piece of information that is sent out on the network. This packet includes the data being sent but also has a header section that contains the source address (where the data is coming from) and the destination address (where the data is headed). If the hacker wants to change who the packet looks like it is coming from, the hacker modifies the source address of the packet.

An example of a spoof attack is the smurf attack. A *smurf attack* is a combination of a denial of service and spoofing. Here's how it works:

1. The hacker pings a large number of systems but modifies the source address of the packet so that the ping request looks like it is coming from a different system.
2. All systems that were pinged reply to the modified source address — an unsuspecting victim.
3. The victim's system (most likely a server) receives so many replies to the ping request that it is overwhelmed with traffic, causing it to be unable to answer any other request from the network.

Eavesdropping attack

An *eavesdropping attack* is when a hacker uses some sort of packet sniffer program that allows him to see all the traffic on the network. Hackers use *packet sniffers* to find out login passwords or to monitor activities. Figure 1-2 shows Microsoft Network Monitor, a program that monitors network traffic by displaying the contents of the packets.

1018 Identifying Types of Attacks

Frame	Time	Src MAC Addr	Dst MAC Addr	Protocol	Description	Src	Obj
1	0.510735	LOCAL	XEROX 000002	Bone	Security Check (0x03)		
2	0.852730	LOCAL	000C29DF8F0C	TCP	...S., len: 0, seq:3110853661-311085366...	XPPE0	
3	8.872759	LOCAL	000C29DF8F0C	TCP	.A..S., len: 0, seq:3238993955-323899395...	10.0.0.	
4	8.882773	LOCAL	000C29DF8F0C	TCP	.A...., len: 0, seq:3110853662-311085366...	XPPE0	
5	8.912816	LOCAL	000C29DF8F0C	LOCAL	FTP	Resp. to Port 1630, '220 Microsoft FTP Serv...	10.0.0.
6	9.022975	LOCAL	000C29DF8F0C	TCP	.A...., len: 0, seq:3110853662-311085366...	XPPE0	
7	10.535149	LOCAL	XEROX 000002	Bone	Security Check (0x03)		
8	12.347756	LOCAL	000C29DF8F0C	FTP	Req. from Port 1630, 'USER administrator'	XPPE0	
9	12.357770	LOCAL	000C29DF8F0C	LOCAL	FTP	Resp. to Port 1630, '331 Password required ...	10.0.0.
10	12.538029	LOCAL	000C29DF8F0C	TCP	.A...., len: 0, seq:3110853662-311085366...	XPPE0	
11	15.492191	LOCAL	000C29DF8F0C	FTP	Req. from Port 1630, 'PASS P@ssw0rd'	XPPE0	
12	15.452220	LOCAL	000C29DF8F0C	LOCAL	FTP	Resp. to Port 1630, '230 User administrator...	10.0.0.
13	15.642493	LOCAL	000C29DF8F0C	TCP	.A...., len: 0, seq:3110853697-311085369...	XPPE0	
14	16.084279	LOCAL	000C29DF8F0C	FTP	Req. from Port 1630, 'PORT 10,0,0,2,6,96'	XPPE0	
15	16.094293	LOCAL	000C29DF8F0C	LOCAL	FTP	Resp. to Port 1630, '200 PORT command succee...	10.0.0.
16	16.094293	LOCAL	000C29DF8F0C	FTP	Req. from Port 1630, 'LIST'	XPPE0	
17	16.904308	LOCAL	000C29DF8F0C	LOCAL	FTP	Resp. to Port 1630, '150 Opening ASCII mode...	10.0.0.
18	16.904308	LOCAL	000C29DF8F0C	TCP	...S., len: 0, seq:3741928937-374192893...	10.0.0.	
19	16.904308	LOCAL	000C29DF8F0C	TCP	.A..S., len: 0, seq:3112918220-311291822...	XPPE0	
20	16.914322	LOCAL	000C29DF8F0C	LOCAL	TCP	.A...., len: 0, seq:3741928938-374192893...	10.0.0.
21	16.914322	LOCAL	000C29DF8F0C	LOCAL	FTP	Data Transfer To Client, Port = 1632, size 13	10.0.0.
22	16.914322	LOCAL	000C29DF8F0C	LOCAL	TCP	.A..F, len: 0, seq:3741928951-374192895...	10.0.0.
23	16.914322	LOCAL	000C29DF8F0C	TCP	.A...., len: 0, seq:3112918221-311291822...	XPPE0	
24	16.944365	LOCAL	000C29DF8F0C	TCP	.A..F, len: 0, seq:3112918221-311291822...	XPPE0	
25	16.954380	LOCAL	000C29DF8F0C	LOCAL	TCP	.A...., len: 0, seq:3741928952-374192895...	10.0.0.
26	17.054824	LOCAL	000C29DF8F0C	TCP	.A...., len: 0, seq:3110853723-311085372...	XPPE0	
27	17.054824	LOCAL	000C29DF8F0C	LOCAL	FTP	Resp. to Port 1630, '226 Transfer complete.'	10.0.0.
28	17.254812	LOCAL	000C29DF8F0C	TCP	.A...., len: 0, seq:3110853723-311085372...	XPPE0	
29	20.559564	LOCAL	XEROX 000002	Bone	Security Check (0x03)		

Figure 1-2:
Using Network Monitor to analyze FTP logon traffic.

Notice in Figure 1-2 that the highlighted packet (frame 8) shows someone logging on with a username of administrator; in frame 11, you can see that this user has typed the password P@ssw0rd. In this example, the hacker now has the username and password of a network account by eavesdropping on the conversation!

Man-in-the middle

While an eavesdropping attack involves the hacker reviewing information, a *man-in-the-middle attack* involves the hacker monitoring network traffic but also intercepting the data, modifying the data, and then sending the modified result out. The person the packet is destined for never knows that the data was altered in transit.

Session hijacking

A *session hijack* is similar to a man-in-the-middle attack, but instead of the hacker intercepting the data, altering it, and sending it to whomever it was destined for, the hacker simply hijacks the conversation, known as a *session*, and then impersonates one of the parties. The other party has no idea he is not communicating with the original partner.

Ensure that you are familiar with the different types of network-based attacks for the A+ exams.



Software-based attacks

Just as there are a number of different types of network attacks, there are a number of software attacks as well. A *software attack* is an attack through software that a user runs. The most popular software attacks are mentioned in the sections that follow.

Trojan horse

A *Trojan horse* is a piece of software that a user is typically tricked into running on the system, and when the software runs, it does something totally different than what the user expected it to do. For example, a typical Trojan horse attack is with a program called NetBus. NetBus is an example of a Trojan horse program that is sent as a file called `patch.exe`. The user receiving the file, typically through an e-mail, believes that the file will fix a security issue. The problem is that `patch.exe` is a Trojan horse, and when that horse starts running, it opens the computer up to allow a hacker to connect to the system.

The hacker then uses a client program, like the one shown in Figure 1-3, to connect to the system and start messing with the computer. The hacker can do things like launch other programs, flip your screen upside-down, eject your CD-ROM tray, watch your activity, and modify or delete files!

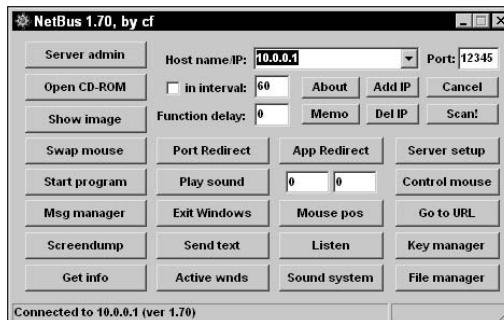


Figure 1-3:
Using
NetBus to
control a
user's
computer.

Virus

A *virus* is a program that causes harm to your system. Typically, viruses are spread through e-mails and are included in attachments such as word processing documents and spreadsheets. The virus can do any of a number of things — it can delete files from your system, modify the system configuration, or e-mail all your contacts in your e-mail software. To prevent viruses, you should install antivirus software and not open any file attachments that arrive in your e-mail that you are not expecting.

Worm

A *worm* is a virus that does not need to be activated by someone opening the file. The worm is *self-replicating*, meaning that it spreads itself from system to system, infecting each computer. To protect against a worm, you should install a firewall. A *firewall* is a piece of software or hardware that prevents someone from entering your system.

Logic bomb

A *logic bomb* is malicious software that could run every day, but the software was designed to wreak havoc on your system on a certain date and time. The scary thing about logic bombs is that they seem like useful software until the day the programmer decides it will become malicious!

Understanding Physical Security

You should implement security in many places. One of the most overlooked areas is physical security. *Physical security* has nothing to do with software; it refers to how you protect your environment and systems by making sure that a person cannot physically access the system. For example, many companies use a numeric keypad to secure the entrance to a facility. In order to get into the facility, you must enter a valid combination on the keypad in order to open the door. Figure 1-4 shows a numeric keypad lock used to enforce physical security.

Another example of physical security is the server room. Most server room doors are locked with a numeric padlock or a key. In order to get access to the server room, you need the key or the correct number for the keypad. Higher-security server rooms sometimes even require fingerprint or retinal scans from anyone trying to enter the room.

The benefit of locking your servers in the server room is that hackers cannot boot off a bootable CD-ROM, which could bypass the operating system entirely. After they have bypassed the operating system, they typically can bypass a lot of the security because they have booted to a totally different operating system.



You can apply security best practices like the ones used by companies with their servers to your home systems. For example, to help secure your home system you may want to prevent booting from a CD-ROM so that an unauthorized person cannot try to bypass your Windows security.



Figure 1-4:
A numeric keypad
used to enforce
physical
security.

In order to protect your systems, follow these physical security best practices:

- ◆ **Server placement:** Lock your servers in a room that only a select few individuals have the key for.
- ◆ **Disable boot devices:** You can help secure the systems by disabling the ability to boot from a floppy disk or CD-ROM in the CMOS setup on the systems.
- ◆ **Set CMOS password:** Because most hackers know how to go to CMOS and enable booting from CD-ROM, you want to make sure that you set a password on CMOS so that a hacker cannot modify your CMOS settings. Figure 1-5 shows a CMOS password being enabled.



Check out Book II, Chapter 4, to get the lowdown on reconfiguring your CMOS settings.

- ◆ **Disable network ports:** To ensure that a hacker doesn't enter your office, plug into the network, and then start performing a number of network attacks, ensure that network ports in lobbies and front entrances are disabled unless an administrator enables them.

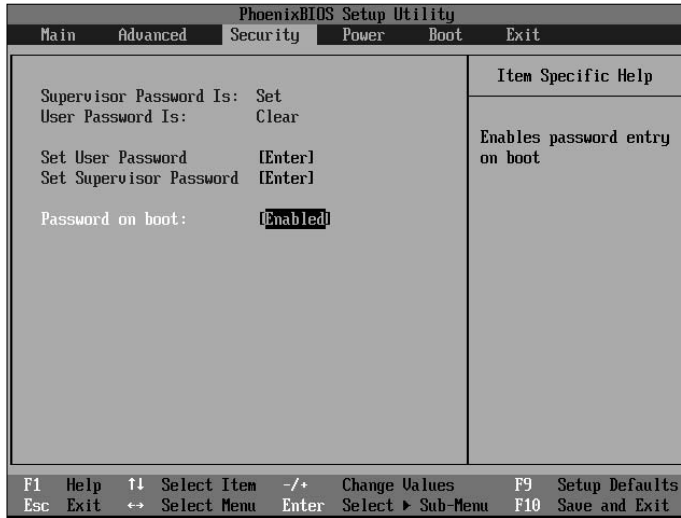


Figure 1-5:
Enabling
the CMOS
password.

- ◆ **Lockdown cable:** A lockdown cable is a cable that you connect to laptops, projectors, and other types of office equipment that locks the device to a table or desk — unless unlocked. Figure 1-6 shows a lockdown cable being used to secure a laptop. A lockdown cable usually connects to a hole in the side of the computer equipment that usually has a picture of a lock next to it.



Figure 1-6:
A lockdown
cable is
used to
secure
computer
equipment
to a desk.



Remembering ways to physically secure your systems will help you with the security portion of the A+ exam. Be sure to place critical systems in locked rooms and lock down equipment that is accessible by the public.

Understanding Authentication and Authorization

After you have physically secured your environment, you then want to focus on the people who access your systems and network. The next step after implementing physical security is to ensure that persons who have entered your server room or have a connection to a network port are authorized to log on to the network. Logging onto the network is known as *authentication*.

Authentication

Authentication is the process of proving one's identity to the network environment. Typically, authentication involves typing a username and password on a system before you are granted access, but you could also use biometrics to be authenticated. *Biometrics* are the use of one's unique physical characteristics, such as a fingerprint or the blood vessels in one's retina, to prove one's identity. Figure 1-7 shows a fingerprint reader that is used to scan your fingerprint in order to log on.




Figure 1-7: A fingerprint reader is an example of biometrics used for authentication.

Here's a quick look at what happens when you log on to your system with a username and password. When you type a username and password to log on to a system, that username and password are verified against a database, known as the *user account database*, which has a list of the usernames and passwords that are allowed to access the system. If the username and password you type are in the user account database, you are allowed to access the system — otherwise, you get an error message and aren't allowed to access the system.

The name of the account database that stores the usernames and passwords is different depending on the environment. In a Microsoft network, the account database is known as the *Active Directory Database* and resides on a server known as a *domain controller* (shown in Figure 1-8).

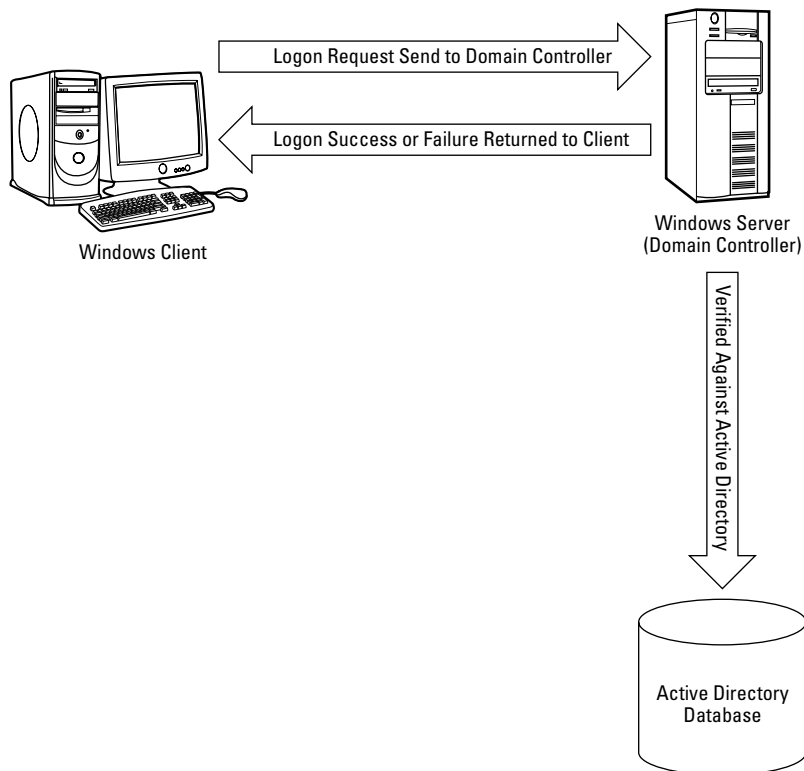


Figure 1-8: Logging on to Active Directory in a Microsoft network environment.

Generating the access token

When you log on to a Microsoft network environment, the username and password you type are placed in a logon request message that is sent to the domain controller to be verified against the Active Directory Database. If the

username and password that you have typed are correct, then an access token is generated for you. An *access token* is a piece of information that identifies you and is associated with everything you do on the computer and network. The access token contains your user account information and any groups you are a member of. When you try to access a resource on the network, the user account and group membership in the access token are compared against the permission list of a resource. If the user account in the access token or one of the groups contained in the access token are also contained in the permission list, then you are granted access to the resource — if not, you get an access-denied message.

If you don't have a server-based network environment and you are simply running Windows 2000 Professional or Windows XP, when you log on, the logon request is sent to the local computer — to an account database known as the *Security Accounts Manager (SAM)* database. When you log on to the SAM database, an access token is generated as well, and that helps the system determine what files you can access.

Smart card

Another type of logon supported by network environments today is the use of a smart card. A *smart card* is a small, ATM card-like device that contains your account information. You insert the smart card into a smart card reader that is connected to a computer, and then you enter the PIN (Personal Identification Number) associated with the smart card. This is an example of securing an environment by forcing someone to not only have the card but also know the PIN.

Strong passwords

It's really hard to talk about authentication without talking about ensuring that users create strong passwords. A *strong password* is a password that is very difficult for hackers to guess or crack because it contains a mix of upper- and lowercase characters, contains a mix of numbers and letters, and is a minimum of six characters long.

Authorization

After a user has logged on and an access token is created, the user may start trying to access resources such as files and printers. In order to access a file, folder, or printer on the network, the user must be authorized to access the resource. *Authorization* is the process of giving a user permission to access a resource. Do not confuse authentication and authorization — you must be first authenticated to the network, and once authenticated, you can then access the resources you have been authorized for.

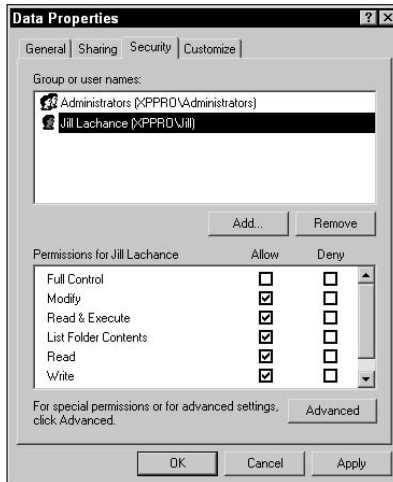
Using Strong Passwords

A number of years ago, I had a coworker who was always trying to get me to guess his passwords. He thought I had some magical trick or program that was cracking them, but all I was doing was guessing his passwords. I remember one time he changed it and I couldn't guess it, until one night we were at a social function for work and all he talked about were the Flyers hockey team. I remember sitting there thinking, "I bet that's his password." Sure enough, the

next day at work, I tried `flyers` as his password, and it worked. Now the lesson here is that he should have at least mixed the case of the word `flyers` to make something like `fLYeRs`, or even better, thrown a symbol in there by replacing the "s" with a "\$." I would have had a much harder time trying to guess his password if he had used `fLYeR$` instead. This is an example of a strong password.

In order to authorize access to a resource, you set permissions on the resource. For example, if you want to allow Jill to access the accounting folder, you need to give Jill permission to the accounting folder, as shown in Figure 1-9.

Figure 1-9: Using permissions to authorize which users are allowed to access the resource.



In Figure 1-9, you can see that the Administrators and Jill have access to the resource. No one else is authorized to access the resource. You find out how to set permissions in the next chapter, but for now, make sure you understand the difference between authentication and authorization.

Methods of Securing Transmissions

After you have authenticated users and authorized them to access certain parts of the network, you should then consider methods of securing information while it travels along the network cable.

Most network communication is sent along the network wire in *cleartext*, meaning that anyone connected to your network can read the information. But if the information is traveling across the Internet, anyone can view that information if it is passed in cleartext.

Most Internet protocols, such as HTTP, send information in cleartext, and it is up to the people who set up the servers that use these Internet protocols to encrypt the information before it is released to the Internet. *Encrypting* the information means that the information is run through a mathematical calculation that generates an altered version of the information, known as a *result*. For example, the words “Glen Clarke” could be encrypted to look like “7y3i s3fk4r.” This is encrypted information, and if anyone intercepts and views it when it is traveling across the wire, the information means nothing to him or her.

A great example is if you were to type your credit card number into a Web site. You don’t want that credit card number to be viewed as you send it from your client computer to the server, so be sure that the Web site you enter the credit card number into is encrypting the traffic. You can tell by the lock icon that appears in the Web browser, as shown in Figure 1-10.

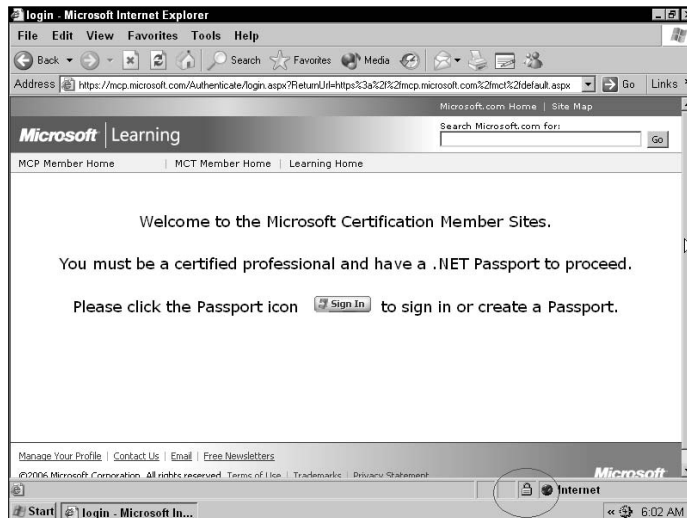


Figure 1-10: Identifying a secure site by locating the lock in Internet Explorer.

1028 *Don't Forget about Data Protection*



It is important for the A+ exam that you understand popular methods of encrypting traffic. You can use a number of technologies, such as

- ◆ **Secure Sockets Layer (SSL):** SSL is a protocol that is used to encrypt different types of Internet traffic. For example, you could use SSL to encrypt HTTP traffic by applying a digital certificate to the Web site. The *digital certificate* contains the key that is used to encrypt and decrypt the traffic.
- ◆ **Internet Protocol Security (IPSec):** IPSec is a protocol that can encrypt all TCP/IP traffic between systems. As a network administrator, you configure IPSec on the server and the clients with the same key, which is used to encrypt and decrypt network traffic. Due to the configuration, it is an unlikely solution for a Web site but is a great way to encrypt traffic on your network.
- ◆ **Virtual Private Network (VPN):** A VPN allows a user to connect across the Internet to a remote network, typically her office network, and send information between her system and the office network securely. The information is secured because the VPN technology used creates an encrypted tunnel between the user and the office network — any data that travels through the tunnel is encrypted.

The preceding sections touch on a number of places that require security. Here's a quick overview of the security steps I've discussed so far:

- ◆ You should secure your office environment first from physical access by unauthorized persons.
- ◆ You should set up a system for authentication, which is the idea that users must log on to the network.
- ◆ After users log on to the network, they must be authorized to access resources.
- ◆ When you allow someone to access resources, make sure that you encrypt the traffic while it is in transit, especially if the information is transmitted outside your own network.

Don't Forget about Data Protection

In this section, you find out about how to secure your data environment from a hacker or malicious user. When securing your systems, you want to protect the systems from a person who damages information or systems with or without intent. You want to be sure to secure your environment from hackers, but at the same time, you want to protect your systems from users on the network who may cause damage without meaning to. Accidents can

always happen, so be sure to prevent accidents from happening by following the best practices in the following sections.

Destroying data

Most office environments have strict policies in place to help secure confidential information. Shredding paper documents that contain personal or confidential information is a no-brainer, and computerized data should be no different. It is important for the company to have strict guidelines on how to destroy data that resides on computer hard drives. Destroying data that resides on a computer hard disk typically involves shredding the computer hard drive with a huge shredding machine — or destroying the drive another way such as sanding the platters down to nothing.

I have talked to some customers who used to destroy drives by driving spikes through them, but what they found was that the data around the hole that the spike put in the drive could still be read! These customers now disintegrate the drive in a huge “shredder,” while other customers sand the drives right down to nothing. Either way, if securing the data is a concern, make sure to physically destroy the disk that contains the data.

Instead of destroying the drives, some companies use a shredder *application* that writes a bunch of 1s to the drive, thereby overwriting the previous data. These applications typically overwrite the drive a number of times because hackers can retrieve the data from disk even after it has been overwritten a few times. If you are purchasing shredding software, be sure to investigate how many overwrite operations the software performs. I recommend using software that overwrites at least 7 times.

Backing up data

A big part of securing the data environment is not only setting the permissions but also ensuring that you create a good backup and restore strategy. It is important to identify which files are critical to the operation of the business and should be backed up. You also want to be familiar with all types of information used by your company. For example, you may depend on e-mail, so you want to make sure you back up your e-mail server along with any files that exist in shared folders. If your company stores important data in databases, you want to make sure that you back up those databases as well.

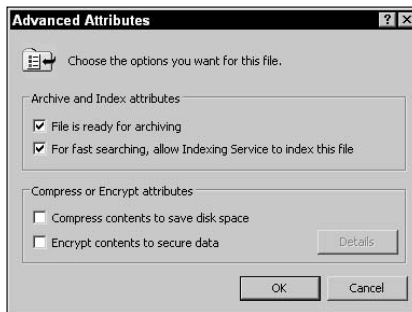
Backup review

You can find out more about backups in Book VII, Chapter 3, but for the exam here are some of the key points you need to remember:

When you perform a backup, the operating system keeps track of which files have been changed since the last backup by setting the *archive bit*. The archive bit is an attribute of the file that tells the system that the file has

changed. To view the archive bit within Windows XP, right-click the file and choose Properties. In the Properties dialog box, click the Advanced button — the Advanced Attributes dialog box appears (shown in Figure 1-11).

Figure 1-11:
Viewing the
archive bit
in Windows
XP.



The first option, File Is Ready for Archiving, is the archive bit. When this check box is checked it means that the file needs to be backed up because it has changed.

Before you perform a backup, you must first decide what type of backup to perform. Each backup type deals with the archive bit a little differently. There are three major types of backup, which are discussed below:

- ◆ **Full backup:** A full backup copies any files that you select, whether the archive bit is set or not, and clears the archive bit on any file that is backed up — essentially recording the fact that the file has been backed up.
- ◆ **Differential backup:** A differential backup copies any files that have changed, but it doesn't clear the archive bit; thus, there is no record that the files have been backed up. The benefit is that the next time you do the backup, the files will be backed up again because the archive bit has not been cleared. As far as the operating system is concerned, the file has not been backed up since it was changed.
- ◆ **Incremental:** An incremental backup copies any file that has changed and then clears the archive bit on any files that are backed up. So if a file is copied during an incremental backup, because the backup process clears the archive bit, the file won't be backed up during subsequent incremental backups unless the file changes again.



For the exam, be familiar with the difference between a full backup, incremental backup, and differential backup. Also know which backup types clear the archive bit.

Tape rotation and offsite storage

You want to make sure that you take the time to rotate tapes so that the same tape is not being used all the time. You also want to make sure that you store a backup offsite in case of a disaster such as flood or fire. It is important that you are able to recover the system no matter what happens.

Test restore operations

As a last point with backup strategy best practices, you want to ensure that you test restorations frequently to ensure that you can recover information from backup without any problem. You don't want to find out that the backups are bad when management is hanging over your shoulder waiting for the company network to come back online! Be sure to perform regular test restorations.

Implementing RAID solutions

To help secure your data, you not only want to make sure you have good backups, but you also want to ensure that you are implementing some form of a *RAID* solution. *RAID (Redundant Array of Inexpensive Disks)* is covered in detail in Book II, Chapter 5, so in this section, I review the different types of RAID volumes supported in Windows 2000 Server and Windows Server 2003 and ensure that you understand that RAID solutions are a way of helping secure data.

RAID is a way of storing duplicated data on multiple disks so that if one disk goes down, the data is still available to the users because other disks in the RAID array have a copy of the data. The benefit of RAID over backups is that with the RAID solution, the user never knows that a drive has failed because the other drive is supplying all the data. You still need the backups in case both drives fail or some disaster happens, like a flood or fire — destroying the system and all of its drives.

There are a number of different types of RAID solutions. The ones provided by the Windows Server operating systems are as follows:

- ◆ **RAID Level 0:** Also known as a *striped volume* in Windows, RAID Level 0 writes different parts of the data to different disks at the same time. The benefit of a striped volume is that you get a performance benefit by writing the data at the same time to two different disks, essentially taking less time to read or write to the file. Note that the data is split between both drives, and there is no duplication — which means that this is not really a redundant solution.
- ◆ **RAID Level 1:** Also known as a *mirrored volume* in Windows. A mirrored volume duplicates the data stored on one disk to another disk. If one disk fails, then the other disk has a copy of the data.

- ◆ **RAID Level 5:** Also known as a *RAID 5 volume* in Windows. A RAID 5 volume requires a minimum of three drives and writes to all drives in the solution like a striped volume. A RAID 5 volume is different than a striped volume in the sense that it does store redundant data, known as parity data, on one of the disks. The redundant data is used to calculate the missing data when a disk goes missing, ensuring that users can still retrieve the data without noticing a problem.



Ensure that you are comfortable with the RAID levels when preparing for the exam. Check out Book II, Chapter 5, to learn how to create volumes in Windows 2000, XP, and Server 2003.

Getting an A+

This chapter introduces you to a number of security-related terms that you need to understand before taking your first A+ exam. The following are some key points to remember when preparing for the exam:

- ◆ *Authentication* is the process of proving your identity to the network, while *authorization* is the process of determining whether you are allowed to access a resource or not after you have been authenticated.
- ◆ Hackers take many different approaches to compromise a system. You should ensure that you protect your environment from both network-based and software-based attacks and that physical security is in place.
- ◆ A *denial of service* (DoS) is an attack on a system or network that prevents the system or network from performing its regular function.
- ◆ *Social engineering* is a popular type of attack that involves the hacker compromising security by tricking an employee through social contact. The social engineer might entice the user to divulge confidential information or may trick the user into running a program that does harm to the system.
- ◆ You secure network traffic by *encrypting* traffic between two systems by using technologies such as SSL and IPSec. Administrators typically use SSL to encrypt Web traffic and IPSec to encrypt internal or VPN traffic.
- ◆ Securing your data involves not only protecting resources with permissions but also protecting your data by following proper data destruction procedures and backup strategies and creating redundant disk solutions.

Prep Test

- 1** What type of attack involves the hacker tricking a user through social contact?
 - A Password attack
 - B Eavesdrop attack
 - C Man-in-the middle attack
 - D Social engineering attack

- 2** What type of attack involves the hacker using a packet sniffer and trying to view confidential information traveling over the network?
 - A Password attack
 - B Eavesdrop attack
 - C Man-in-the-middle attack
 - D Social engineering attack

- 3** What type of attack involves the hacker causing your system or network to become unresponsive to valid requests?
 - A DoS attack
 - B Eavesdrop attack
 - C Man-in-the-middle attack
 - D Password attack

- 4** What type of RAID volume duplicates the data fully on two disks?
 - A Striped volume
 - B Mirrored volume
 - C RAID 5 volume
 - D RAID Level 0

- 5** What type of attack involves the hacker capturing network traffic, altering the data, and sending it on to its destination?
 - A Password attack
 - B Eavesdrop attack
 - C Man-in-the-middle attack
 - D Social engineering attack

- 6** What type of software-based attack involves a program performing an unexpected function at a certain date or time?
- A Virus
 - B Worm
 - C DoS
 - D Logic bomb
- 7** You wish to ensure that a hacker cannot boot off a CD-ROM or floppy disk to bypass the operating system; what should you do?
- A Set a password in Windows.
 - B Disable the CD-ROM and floppy as bootable devices in CMOS. Also put a password on CMOS.
 - C Disconnect the CD-ROM and floppy drive.
 - D Ensure that the CD-ROM and floppy are first in the boot order — before the hard disk.
- 8** What type of RAID volume stripes the data across all disks but also contains parity information used to recalculate missing data?
- A Striped volume
 - B Mirrored volume
 - C RAID 5 volume
 - D RAID Level 0
- 9** What should you purchase for each laptop to help protect it from theft?
- A Flash drive
 - B Driver disk
 - C A Doberman pinscher
 - D Lockdown cable
- 10** What type of authentication device requires you to know the PIN when you place the card into the reader?
- A Fingerprint scanner
 - B Smart card
 - C Biometric device
 - D Retinal scanner
- 11** Which of the following are forms of biometrics? (Select all that apply.)
- A Fingerprint scan
 - B Smart card
 - C Username and password
 - D Retinal scan

- 12** What type of software-based attack involves a virus spreading itself from system to system without needing to be activated by a user?
- A Trojan
 - B Worm
 - C DoS
 - D Logic bomb
- 13** What type of backup copies the files that have changed and does not clear the archive bit?
- A Full backup
 - B Incremental backup
 - C Differential backup
 - D Copy
- 14** What technology is typically used to encrypt traffic between a Web server and Web browser?
- A DoS
 - B IPSec
 - C Smart card
 - D SSL
- 15** In high-security environments, what should you do with old hard drives?
- A Donate them to charity
 - B Recycle them
 - C Physically destroy them
 - D Drive a spike through them
- 16** Which of the following is the strongest password?
- A thisisalongpassword
 - B P@ssw8rd
 - C password
 - D StrongPassword
- 17** Where should your company's servers be located?
- A At the front door
 - B In a manager's office
 - C At the reception desk
 - D In a locked room

- 18** What technology is used to encrypt all TCP/IP traffic?
- A DOS
 - B IPSec
 - C Smart card
 - D SSL
- 19** What type of attack involves the hacker modifying the source address of the packet?
- A Spoof attack
 - B Eavesdrop attack
 - C Man-in-the middle attack
 - D Social engineering attack
- 20** When is an access token generated?
- A Only during logon
 - B During logon and automatically every 90 minutes
 - C When the system is turned on
 - D When the system is turned on and automatically every 90 minutes

Answers

- 1** D. Social engineering is a type of hack that involves contacting victims through phone or e-mail and tricking them into doing something that compromises company security. *See "Social engineering attacks."*
- 2** B. An eavesdropping attack is when a hacker monitors network traffic to try to capture information that could be useful in another attack. *Review "Eavesdropping attack."*
- 3** A. A denial of service (DoS) attack is when a hacker consumes all of the system's processing power or bandwidth so that it is unable to perform its normal job. *Check out "Denial of service."*
- 4** B. A mirrored volume is used to create a full duplicate of the data on two different disks. *Peruse "Implementing RAID solutions."*
- 5** C. A man-in-the-middle attack is when the hacker captures data traveling on the wire, alters the data, and sends it to the person it was originally destined for. The parties involved have no idea they are dealing with altered information. *Take a look at "Man-in-the-middle."*
- 6** D. A logic bomb is placed in the code of the program so that on a certain date, the program harms the system. Before that triggered date, the program runs normally and offers benefits to the user. *Peek at "Logic bomb."*
- 7** B. Ensure that a hacker cannot boot off a CD-ROM or floppy disk by disabling them as bootable devices in CMOS. Also be sure to set a CMOS password so that hackers cannot easily enter CMOS and change the boot settings. *Look over "Understanding Physical Security."*
- 8** C. A RAID 5 volume spreads the data across a number of disks and then calculates parity information, which is used to generate missing data when a disk fails. *Study "Implementing RAID solutions."*
- 9** D. A lockdown cable is used to secure the laptop to a desk to help prevent the laptop from being stolen. *Refer to "Understanding Physical Security."*
- 10** B. A smart card is a small credit card–like authentication device that is inserted into a smart card reader. After the card is inserted into the reader, the user then types a PIN to be authenticated. *Examine "Smart card."*
- 11** A, D. Biometric devices involve authenticating a user through the user's unique physical characteristics. Fingerprint scans and retinal scans are popular biometric authentication methods. *See "Authentication."*
- 12** B. A worm is a self-replicating virus that bounces from system to system. You can use a firewall to stop the worm from replicating to your system. *Review "Worm."*

Fundamentals of Security

- 13** C. A differential backup only backs up the files that have changed since the last full backup and then does not clear the archive bit. *Check out "Backup review."*
- 14** D. Secure Socket Layer (SSL) is used to encrypt Web traffic. You can identify whether or not you are on a secure Web site by looking for the lock icon at the bottom of the screen. *Peruse "Methods of Securing Transmissions."*
- 15** C. You want to make sure that you physically destroy the drives if securing data is critical to the business. *Take a look at "Destroying data."*
- 16** B. P@ssw8rd is the strongest password listed because it a) is more than six characters long, b) it has a mix of upper- and lowercase characters, and c) it is a mix of numbers, letters, and symbols. *Peek at "Strong passwords."*
- 17** D. You should keep your servers in a locked room. You want to make sure that you physically secure your systems, along with all the other security measures. *Look over "Understanding Physical Security."*
- 18** B. Internet Protocol Security (IPSec) is used to secure all TCP/IP traffic. *Study "Methods of Securing Transmissions."*
- 19** A. A spoof attack is when the hacker modifies the source address, trying to hide the origin of the packet. *Refer to "Spoofing."*
- 20** A. The access token is generated only during logon. *Examine "Generating the access token."*

Chapter 2: Implementing Security

Exam Objectives

- ✓ **Implementing users and groups**
- ✓ **Configuring a security policy**
- ✓ **Assigning permissions**
- ✓ **Implementing firewalls**

In this chapter, you find out how to implement security best practices on systems at home or at the workplace. The preceding chapter introduces terms such as *authentication* and *authorization*; this chapter demonstrates how to perform such tasks. You find out how to create a user account that can be used for authentication and how to authorize the user to access a folder or perform an action within the operating system. This chapter will ensure that you know how to perform basic security-related tasks!

It is important to understand when thinking about network security that security is to be implemented at multiple layers, meaning that you cannot focus on just one security-related feature — you want to implement multiple security features to secure your environment. For example, a number of people feel that their systems are secure because they have a firewall. They don't realize that the firewall protects the system only from attacks coming across the network. What if the hacker is in the same room as the computer? The firewall is of no use at that point, so you need to ensure that you implement other security features to protect the system from all potential threats.

Securing Systems through BIOS

When securing systems, your first security concern is physical access. This involves ensuring that critical systems, such as servers, are in locked rooms that are not accessible to unauthorized users. Physically securing systems could also involve changing some of the CMOS settings, such as boot device order, power-on password, and CMOS password.

Changing these settings in CMOS is different for each type of system, but the first thing you have to do is enter CMOS. Normally, you press Delete, F1, F2, or F10 when the system is booting.

After the system is booted, you will find the following settings in the CMOS setup program to help secure the system:

- ◆ **BIOS Password:** Usually found in the security section of CMOS, you can set a power-on password (also known as a user password), which is a password that anyone who wants to use the system must type. You may also set an admin password, which is a password that must be known by anyone who wants to change CMOS settings.
- ◆ **Boot Devices:** In CMOS, you can control what devices the computer can boot from. Most computers today can boot from CD-ROM, floppy disk, hard disk, network, and USB removable drives. It is important to understand that if you allow a computer to boot from CD-ROM, it is possible that a hacker can boot from a CD and bypass all security enforced by your operating system.
- ◆ **Intrusion Detection:** Most systems today have an intrusion detection option that will notify you if the computer case has been opened. This is important because instead of stealing the actual computer, a person could take the RAM or hard drive out of the computer, which is easier to hide in a duffel bag. Make sure that the intrusion detection option is enabled, and also be sure to lock the computer cases so they cannot be removed easily.

Implementing Users and Groups

In this section, you find out how to create user accounts that can be used to log on to the system and how to create groups to organize users together as a single object that permissions can be assigned to.

Creating user accounts

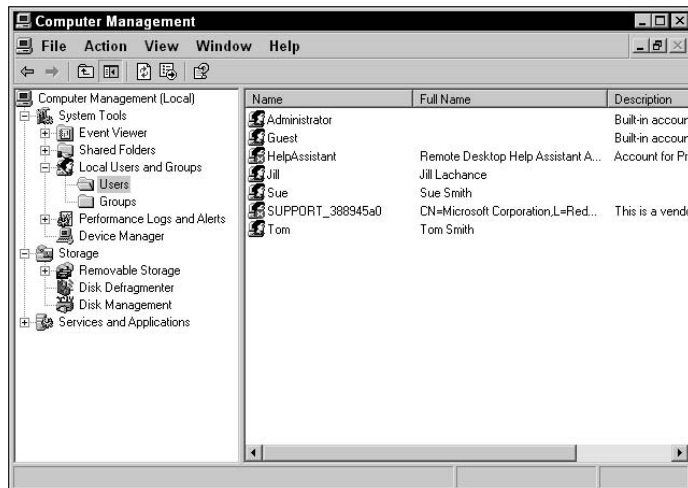
To secure the Windows operating system from unauthorized access, you can create a user account for each person who is allowed to use the system. Anyone who doesn't have a user account will be unable to log on to the system and, as a result, will not be able to use the computer. The other benefit of creating user accounts is that even if a person has a user account and logs onto the system, he or she may not be able to access a file because you have not given permission to that user to access the file.

To create a user account in Windows 2000 or Windows XP systems, you use the Computer Management console. To start the Computer Management console, right-click My Computer and choose Manage.

In the Computer Management console, expand Local Users and Groups and select the Users folder (shown in Figure 2-1). In the Users folder, you will

notice some user accounts on the right side. These user accounts are *built-in accounts*, meaning that they were built by the operating system or by a piece of software you have installed.

Figure 2-1:
Creating
user
accounts
and groups
in the
Computer
Manage-
ment
console.



Two built-in accounts you should be familiar with for the A+ exam are:

- ◆ **Administrator:** The *administrator account* is the built-in account in Windows that has full access to the system and can manage all aspects of the computer. During the installation of Windows 2000/XP/2003, you were asked what you wanted to set as the password for the administrator account; you use that password to log on with the username of *administrator*. When you do log on as administrator, you can change any settings on the system. A normal user account cannot change major settings on the system such as the time, installing software, or any changes that affect the system. To make these types of changes you need to log on as administrator to make changes.
- ◆ **Guest:** Users can use the *guest account* if they don't have an actual user account. When they try to access the system, they are authenticated as *guest*. The guest user inherits any permissions the guest account has on the system. There is one hook to this scenario — by default, the guest account is disabled, meaning that it is not available for use. Due to the security concerns of not requiring someone to log on, Microsoft has disabled the account. A disabled account appears with a red "X" on it and cannot be used.



For the exam, remember that there are two default accounts built in Windows — the administrator account and the guest account. The administrator account has full access to the system while the guest account is used for temporary access to the system. Also note that the guest account is disabled by default.

Now that you have identified the two major built-in accounts, you can create your own user accounts. To create your own user accounts in the Computer Management console, right-click the Users folder and choose New User. The New User dialog box appears (shown in Figure 2-2). Fill in the following account details:

- ◆ **User Name:** This will be the name that the user will use to log on to the system. Typically, it is a short version of the full name. For example, my full name is Glen Clarke so I might use `gclarke` as my username. A username is also known as the logon name.
- ◆ **Full Name:** This is typically the person's first name and last name. For example, my user account would have `Glen Clarke` as the full name.
- ◆ **Description:** This is a description of the user account. I normally put the person's job role here. For example, if I was an accountant, I might put `Accountant` in the description.
- ◆ **Password:** Type what you want for the user accounts password. The user needs to know this password to log on to the system. Be sure to use good practices with passwords, such as not using words found in the dictionary and using a combination of upper- and lowercase letters, numbers, and symbols. See the preceding chapter for more information about strong passwords.
- ◆ **Confirm Password:** Type the password again in this box. This ensures that you typed what you thought you typed.
- ◆ **User Must Change Password at Next Logon:** Set this option if you want to force the user to change the password the first time he logs on. This ensures that you don't know the user's password because the password you originally set is overwritten.
- ◆ **User Cannot Change Password:** Set this option if you don't want the user to be able to change the password. This ensures that the password you set is the password the user is using.
- ◆ **Password Never Expires:** In a password policy, you can specify that passwords must be changed every so many days. That policy applies to all users except for any accounts that have Password Never Expires activated. You might use this option if you have two employees sharing a user account.
- ◆ **Account Is Disabled:** If you want to disable an account at any time, you can set this option. A disabled account is unusable until you enabled it again.

Figure 2-2:
Creating
a user
account in
Windows
XP.

The screenshot shows the 'New User' dialog box in Windows XP. The fields are filled with the following information:

- User name: bsmith
- Full name: Bob Smith
- Description: Accountant
- Password: [masked]
- Confirm password: [masked]

Below the password fields, there are four unchecked checkboxes:

- User must change password at next logon
- User cannot change password
- Password never expires
- Account is disabled

At the bottom right, there are two buttons: 'Create' and 'Close'.

After you have typed all the account information, click the Create button and then click Close to get rid of the New User dialog box. The user account has been created, and you can start using it right away to log on to Windows.

Creating groups

A *group* in Windows is a collection of user accounts. The benefit of using groups when managing access to resources is that you don't need to assign the same permissions multiple times — you assign the permission to the group, and anyone who is a member of the group receives the permission.

Like user accounts, Windows offers a number of built-in groups. A built-in group has predefined capabilities within Windows. For example, printer operators can manage all printers on the system, and anyone who is a member of the printer operators group will have that capability. The following is a list of some of the popular built-in groups found in Windows 2000/XP/2003:

- ◆ **Administrators:** This group has full access to the system and can change any setting on the system. The administrator account is a member of this group by default, which is why the administrator account is allowed to change any setting on the system.
- ◆ **Backup Operators:** Members of this group are allowed to perform backups and restores on the system.
- ◆ **Account Operators:** Members of this group are allowed to create user accounts. This group is available on the server versions of Windows 2000 and Windows 2003. The benefit of using this group is that if you want someone to be able to manage user accounts, you can place that person in this group instead of in the administrators group and he or she will only be able to manage the user accounts — not the entire system!

- ◆ **Printer Operators:** Members of this group can change any settings on the printers. Essentially, members of this group are trained to troubleshoot the printing environment and then assigned the task of managing all printing problems on the network.
- ◆ **Users:** All user accounts that are created are members of the users group. You can assign permissions to the users group knowing that all users will get the permission.
- ◆ **Power Users:** The power users group is the group on Windows 2000 Professional and Windows XP Professional that allows its users to create user accounts and manage the printing environment. Use this group if the desktop operating system does not have an account operator or a printer operator group.



For the exam be sure to know the default groups that exist in Windows. Some of the more useful built-in groups are account operators, printer operators, and backup operators.

If the built-in groups do not satisfy your needs, you can create your own groups as well. To create your own groups in Windows XP, follow these steps:

1. Click Start and then right-click on My Computer and choose **Manage**. The Computer Management console starts.
2. In Computer Management, expand **Local Users and Groups**.
3. Right-click the **Groups** folder in Local Users and Groups and then choose **New Group**, as shown in Figure 2-3.

Figure 2-3:
Creating a new group in Computer Management with Windows XP.



The New Group dialog box appears.

4. Type the name you wish to use for the group.

In this example, I type Accountants (shown in Figure 2-4).



Figure 2-4:
Filling in
the group
information
for the
accounting
group.

5. Fill in a description for the group in the Description text box.

6. To begin adding members to the group, click the Add button.

The Select Users dialog box appears.

7. Type the name of the user account you want to add and click the Check Name button on the right side.

Windows should underline the account name, indicating that the user account does exist and that you can add it to the group membership.

8. Repeat Step 7 for each account you want to add to the group.

9. After you have added all the accounts to the group, click OK and then click Create to create the group.

Now that you have created the users and placed them into their appropriate groups, you are now ready to assign them permissions.



To practice creating users and groups take a look at Lab 2-1. Lab 2-1 can be found in the Labs .pdf file in the Author directory of the CD-ROM.

Implementing Permissions and Rights

When controlling a user's access to the system, you typically modify the user's *rights* and *permissions*. Microsoft has made a huge distinction between

a permission and a right. A *permission* is a user's level of access to a resource — such as a printer or file — while a *right* is a user's privilege to perform an operating system task. In this section, you discover the difference between permissions and rights within the Windows operating system and how to implement both.

Rights

If you were to log on to your Windows system as just a user account and then double-click the time in the bottom-right corner to change that time, you get an error message indicating that you do not have the privilege to change the time. This is an example of user rights. The user account that you are currently logged in with does not have the right to change the system time, an action that typically has to be performed by an administrative account.

There is a large list of user rights; some of the most popular ones are listed below:

- ◆ **Access this computer from the network:** This right is needed by anyone who wants to connect to the system from across the network. For example, if you wish to connect to a shared folder on Computer A, you need to have the *access this computer from the network* right on Computer A.
- ◆ **Back up files and directories:** This right is needed by anyone who wishes to back up files on the computer. For security reasons, not everyone should be able to perform backups on a system, so Windows controls who can perform a backup via this right.
- ◆ **Change the system time:** In order to change the time on the computer, your user account must be given the *change the system time* right.
- ◆ **Log on locally:** In order to log on to the system by pressing Ctrl+Alt+Delete, you need to have the log-on-locally right. Microsoft classifies a *local logon* as you sitting in front of the computer at the keyboard — a *remote logon* is you connecting from across the network, which is controlled by the first right mentioned in this list.
- ◆ **Shut down the system:** In order to shut down the computer, you must have this right.
- ◆ **Take ownership of files and other objects:** In Windows, the owner of the object, such as a file or folder, always has the ability to change the permissions of the resource. You may want to give selected individuals the take-ownership right so that they can take ownership of a resource and then change the permissions.

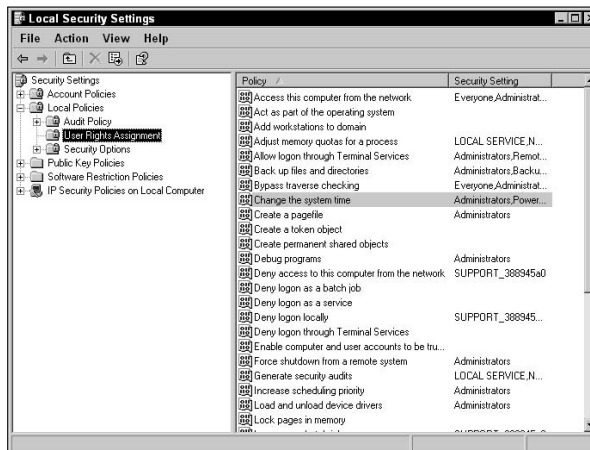
To change the user rights (for example, to assign Bob Smith the right to change the system time), you need to modify the user rights assignments

in the local security policies of the Windows computer. The local security policy controls all security settings for the system. To change the local security policies in Windows XP, follow these steps:

1. Choose **Start** → **Control Panel**.
2. In the Control Panel, choose **Performance and Maintenance** and then **Administrative Tools**, located at the bottom of the window.
3. In the Administrative Tools, double-click **Local Security Policy** to start the Local Security Policy console.
4. To modify the user rights within the local security policy, expand **Local Policies** and then highlight **User Rights Assignments**, as shown in Figure 2-5.

When the User Rights Assignments node on the left side has been selected, you will notice the list of user rights on the right side of the screen in the Details pane.

Figure 2-5: Configuring user rights within Windows allows you to control which actions a user can perform.



5. To modify a user right, double-click the user right.

You will see a list of users or groups that have been assigned that right.

6. To add a user or group to the list, click the **Add User or Group** button and then type the name of the account you wish to add and then click **Check Names** to ensure that Windows recognizes the user account.
7. Click **OK** to add the account to the right you chose (as shown in Figure 2-6) and then click **OK** to close the window.



Figure 2-6: Adding Bob to the change the system time user right.

Permissions

Permissions are different than rights: A right governs an action that can be performed on the computer, but a *permission* is a user's level of access to a resource. For example, you can give a user permission to read or modify a file. Figure 2-7 shows the permissions you can set for a file.

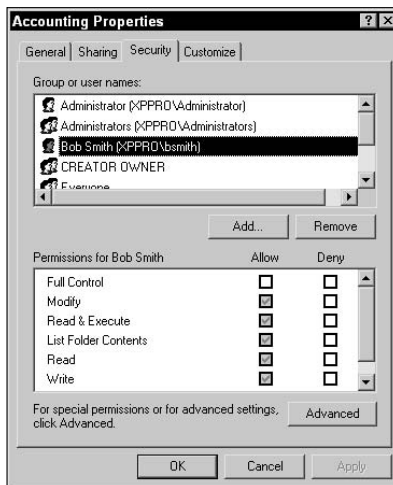


Figure 2-7: Looking at NTFS permissions in Windows XP.



Permissions can be configured only on a partition formatted for NTFS. To obtain an NTFS partition, you can format the partition for NTFS (but lose all existing data), or you can convert the drive to NTFS by using the `convert driveletter: /fs:ntfs` command. When you convert, the existing data on the drive is preserved.

Here are the available permissions:

- ◆ **Read permissions:** What I call the read permission is a combination of the three default permissions — Read, Read and Execute, and List Folder Contents. I personally classify all three as the “read” permission because, at a minimum, this is typically what users need in order to read the file. The Read permission allows you read the contents of a file, the Read and Execute permission allows you to read the contents of the file and execute a program, and the List Folder Contents permission allows you to see the file when you look in the folder.
- ◆ **Modify:** The Modify permission allows a user to read, modify, and delete a file. When given the Modify permission to a folder, a user can also create new files or folders in that folder.
- ◆ **Full Control:** The Full Control permission allows a user to do everything that the Modify permission allows, but the user can also change permissions on the resource or take ownership of the resource. Remember that if someone can take ownership of the resource, that person can change the permissions. The Full Control permission should be used sparingly so that not everyone has the permission to change permissions on you.
- ◆ **Write:** The Write permission is used by the Modify permission to allow users to write to the file or folder. When you choose the Modify permission you will notice that the Write permission is automatically selected.



For the exam, remember that the major difference between the Modify permission and Full Control permission is that full control allows a user to modify permissions and take ownership of the resource on top of being able to modify and delete the resource.

Looking at Figure 2-7, you will notice that a number of existing permissions have gray check boxes next to them. The gray check box means that you are not allowed to change the permission because the permission is being inherited from a parent level. *Permission inheritance* is a feature of Windows that is designed to minimize how much permission management you need to do. With permission inheritance, when you set permission on a folder, that permission applies to all subfolders and files; you don’t need to go to subfolders and files to set the same permission.

When you go to modify the permissions on a folder, however, you need to understand that the existing permissions are being inherited from the parent

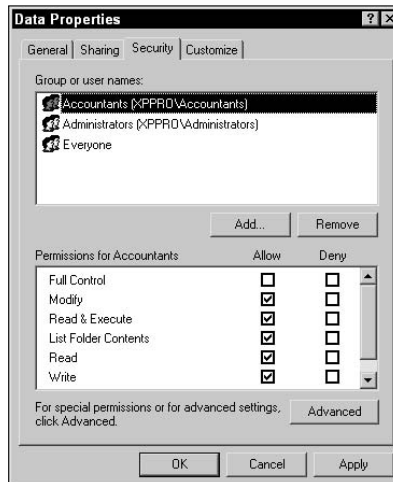


Figure 2-9:
Giving the
accountants
the Modify
permission.

Implementing Auditing

After you have set up security on a Windows system by setting permissions on the folders and files, configuring user rights, and placing users in the appropriate groups, you should make sure that the security of the operating system is effective. To monitor what is happening on the system, you enable auditing. *Auditing* is a feature that notifies you when certain things happen on the system. For example, you may want to be notified if someone fails to log on to the system with a correct username and password because this could be someone trying to guess the password of the account.

To effectively work with the auditing feature in Windows, there are two steps:

- 1. Enable auditing:** You must first enable auditing. To enable auditing, you simply choose what events you wish to audit. The nice thing about auditing in Windows is that you choose which events you care to know about.
- 2. Review the audit log:** After you have enabled auditing, you need to ensure that you monitor the log regularly for any security-related issues. For example, if you notice a failure to log on over and over for the same account, then that is an indication that an account is being hacked.

The following sections offer more details about these two steps.

Enabling auditing

To enable auditing in Windows 2000/XP/2003, modify the Local Security Policy:

1. Choose Start⇒Control Panel.
2. In the Control Panel, choose Performance and Maintenance and then Administrative Tools, located at the bottom of the window.
3. In the Administrative Tools, double-click Local Security Policy to start the Local Security Policy console.
4. In the Local Security Policy console, expand Local Policies and then highlight Audit Policy.

You will notice a list of events that you can enable auditing for on the right side of the screen, called the *Details pane*, which is displayed in Figure 2-10:

- *Audit Account Logon*: Any remote users who are authenticated by this user account database are audited. This is the event to enable auditing on a domain controller.



A *domain controller* is a server in a Microsoft network environment that holds all the user accounts for an entire network. In the corporate world, users log on to the network, not a particular machine, which means that the logon request is sent to the domain controller where the user name and password are checked against a database. The database that holds the user accounts on a domain controller is known as the *active directory* database.

- *Audit Account Management*: Records an event in the log for any user account changes, such as any new accounts that are built, modified, or deleted.
- *Audit Logon Events*: Record the fact that the user logged on from this station, whether or not the account was authenticated from this system.
- *Audit Object Access*: Audits access to a specific folder, file, or printer.



After you enable Object Access Auditing, you need to go to the Security page in the properties of a file, folder, or printer and click the Advanced button. Click the Auditing tab and choose which users and which permissions to audit for. You must perform this step on any folder, file, or printer you wish to audit.

- *Audit Policy Change*: Notification of any change to the security policy.
- *Audit Privilege Use*: Logs when a user takes advantage of any rights you have given that user. For example, if you gave Bob the right to perform backups, you want to know when he actually performs a backup.
- *Audit Process Tracking*: This event will notify you when a process starts or exits.

- **Audit System Events:** Notification of system-related actions, such as restarting or shutting the system down. You may want to be aware when the system is restarted, especially on server operating systems.

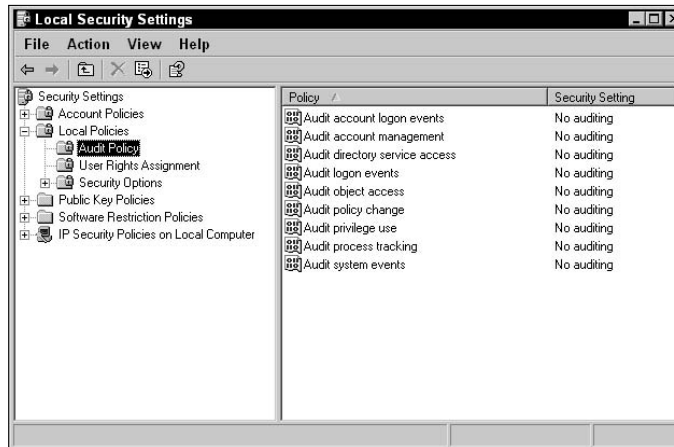


Figure 2-10: Looking at the auditing feature within the Local Security Policy.

5. To enable auditing on one of these events, double-click the event and then choose whether you want to audit the success of that event or the failure.

For example, I probably don't care about the success of logons, so I would choose Failure for that event.


Reviewing the Security Log

After you have enabled auditing on the different events, you then need to view the audited information in the security log of event viewer. To view the audited information in the security log, follow these steps:

1. Choose Start → Control Panel.
2. In the Control Panel, click Performance and Maintenance and then Administrative Tools, located at the bottom of the window.
3. In the Administrative Tools, double-click the Event Viewer to start the Event Viewer console.
4. Select the log that you want to view, and you will notice a number of events on the right side of the screen.

If you select the Security log, as shown in Figure 2-11, any events with a lock are failure events, and any events with a key are successful events. Note in Figure 2-11 that there is an account logon event with a lock indicating a failure to log on.

Figure 2-11: Review the security log that is populated by the auditing feature of Windows.



5. To view a description of a particular event, double-click the event.

Going back to the account logon failure example, you can see the date and time the logon was attempted. You can also view the username that was attempted and the computer that the person used to try to log on to the network.



For the exam remember that after enabling auditing you will review the security events by checking out the security logon event viewer.

Implementing Firewalls

A *firewall* is a piece of software or hardware that is designed to stop information from reaching your system unless you selectively choose certain pieces of information to pass through the firewall. This information is sent in the form of network packets (pieces of data) that are broken down into three parts:

- ◆ **The header** of the packet contains address information, such as source and destination addresses.
- ◆ **The body** of the packet contains the packet data, known as the payload.
- ◆ **The trailer** contains checksum information, which helps ensure that the data has not been tampered with or damaged in transit.

How a firewall works

A firewall is designed to look at the contents of the packet, specifically the header information, to decide whether the data should be allowed into the system or discarded. The firewall uses the source and destination IP addresses from the header, as well as the port number, to help make this decision. A port number represents an application that runs on the system.

For example, the Web server installed on my system runs at my IP address on port 80. The FTP server I am also running on my system uses my IP address but uses port 21 instead of port 80. If I want to allow the public to see my Web site but not my FTP site, then I configure the firewall to allow information to reach port 80, but not port 21. So each TCP/IP application that is running on your system uses a different port number, which is how data is sent to one application and not the other.

My point is that the firewall also uses the port number to decide whether the data should be allowed into your system. For example, I have a Web site at www.gleneclarke.com, so I had to configure my firewall to allow data destined for port 80 to be allowed in. Now, I don't have an FTP server, so I ensured that the firewall disallows data destined for port 21.



It is important to understand that you don't need to open ports on the firewall unless you are hosting your own servers. For example, you don't need to open ports on the firewall to surf the Internet because most firewalls are built to allow responses to data you requested to come back through the firewall.

To enable the firewall feature in Windows XP, follow these steps:

- 1. Go to your network properties by choosing Start⇨Control Panel⇨Network and Internet Connections⇨Network Connections.**
- 2. In the Network Connections window, right-click your LAN connection and choose Properties.**
- 3. In the properties of the LAN connection, click the Advanced tab to view the advanced settings.**

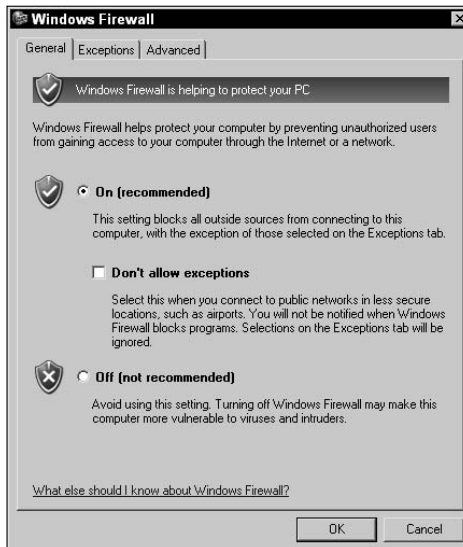
You will notice a Windows Firewall section at the top of the screen.

- 4. Click the Settings button to enable the firewall (as shown in Figure 2-12).**
- 5. Make sure the firewall is set to On.**
- 6. You can also build exceptions for information that is allowed to pass through the firewall by clicking the Exceptions tab.**

On the Exceptions tab, you may select which data is allowed to pass through the network card into the system. You may select an existing application from the list or add a program or port by clicking the Add Program button or Add Port button.

- 7. When you have ensured that the firewall is enabled, click the OK buttons to close the windows.**

Figure 2-12:
Enabling the
Windows
firewall
feature in
Windows
XP.



Creating a DMz

Most companies that want to publish their own Web sites or host other types of servers such as FTP servers or e-mail servers need to allow traffic to reach these types of servers. Placing public servers such as these alongside your private network servers is unrealistic because it means that you need to open the firewall to allow traffic into the network to reach these servers.

Most network administrators create a *demilitarized zone* (DMz) to hold these servers. A DMz is a network segment between two firewalls where you have allowed selected traffic to reach the servers in the DMz. The DMz is different than your private network because you will not allow any content to come into your private network.

Figure 2-13 displays a typical DMz setup. Notice that there are two firewalls, named firewall 1 and firewall 2. Firewall 1 is connecting the DMz to the Internet and will only allow traffic that is destined for the three servers in the DMz to pass through the firewall. The second firewall, named firewall 2, is designed so that no systems from the Internet can pass through it — essentially protecting the private company network from outside access.

Servers that you wish to expose out to the Internet should be placed in a demilitarized zone so that you can selectively choose which type of data is allowed to reach your servers.





Figure 2-13:
Identifying
a DMz.



Implementing Security Best Practices

In the following sections, you discover some basic best practices that can help you secure your environment. These sections are designed to be a summary of features that I discuss throughout the chapters of this book.

Hardening a system

The first thing you can do to secure your system is to *harden* it. Hardening a system means that you remove any software that you are not using and you disable any Windows services that are not needed. The concept of hardening comes from the fact that hackers compromise systems by leveraging software that is installed or running on the system. The less software you have running, the less likely you are to be hacked!

Patching systems

Regularly patching the system by running Windows Update is critical. As Microsoft finds out about security problems with its operating system and software, its programmers fix the problem and deliver the fix through the Windows Update site. To ensure that you are getting the security fixes and patches, you must run Windows Update often. More on this topic in the next chapter.



Microsoft has changed from Windows Update to Microsoft Update so that you can now get updates for more than just the Windows operating system. You can download updates for a number of Microsoft products from the Microsoft Update site, such as Windows and Microsoft Office.

Firewalls

Make sure you turn on the firewall feature in Windows. The firewall helps protect your system from network attacks, but it is not the be-all and end-all of network security. You also need to follow the other best practices presented in this chapter.

Password policies

Stress to your users the importance of using strong passwords. To enforce strong password usage, you can set a password policy in the Local Security Policies.

To set the password policy, follow these steps:

- 1. Choose Start→Control Panel.**
- 2. In the Control Panel, click Performance and Maintenance and then click Administrative Tools, located at the bottom of the window.**
- 3. In the Administrative Tools, double-click Local Security Policy to start the Local Security Policy console.**
- 4. Expand Account Policies and highlight Password Policy.**
- 5. Ensure that users use strong passwords by double-clicking the Password Must Meet Complexity Requirements and choosing Enable to enable the policy (see Figure 2-14).**

This setting ensures that users use passwords with a mix of upper- and lowercase characters, numbers, and symbols and a minimum of six characters. The password will also not contain any part of the username.

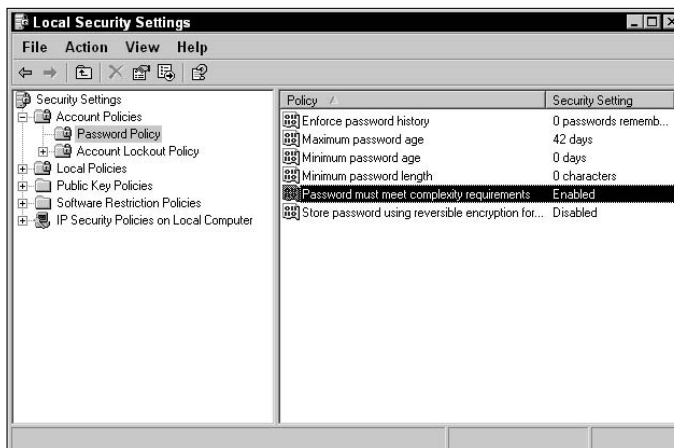


Figure 2-14:
Configuring password complexity in Windows 2000/XP.

Auditing

Make sure that you enable auditing on critical systems so that you will know (hopefully) when the system has been compromised. For example, if a hacker makes his way into the system and builds himself a hidden user account, you will know about it if you have enabled account management auditing.

Use switches instead of hubs

You can enable a number of security features when working with switches instead of hubs on the network. To begin with, switches filter traffic by only sending the data to the port on the switch that the data is destined for. This can add to the security of the network because it is harder for a hacker to monitor network traffic when the port the hacker is using is not getting a copy of all data — just data destined for his system.

The second thing you could do to secure your environment with a switch is disable any unused ports on the switch. This way, if the hacker gets physical access to your network, she cannot simply plug into the switch to get access to the network.

The other thing you could do with more advanced switches is to configure *Virtual Local Area Networks (VLANs)*. A VLAN is a grouping of ports on the switch that are allowed to communicate with one another but cannot communicate with other VLANs on the same switch. For example, I have a 24-port switch that has two VLANs. The first VLAN is made up of the first 12 ports, while the second VLAN is made up of the last 12 ports. Any systems that are plugged into the first 12 ports cannot communicate with the systems that are on the second set of 12 ports, and vice versa. Essentially, you have two networks — but only one switch.

Use anti-virus software

Although we have not touched on anti-virus software, I want to add it as a security best practice — ensure you are using anti-virus software on all of your systems! Anti-virus software is software that is designed to protect your system against viruses. For more information on anti-virus software check out book IX, chapter 3.

Securing wireless

As a last note, I just want to add a few tips here to help secure your wireless environment:

- ◆ **Router password:** After you hook up your wireless router be sure to connect to the router and change the admin password. Most wireless routers ship with no password so be sure to protect your router by assigning one. Check the documentation that came with your router to find out how to set an admin password.
- ◆ **Set the SSID:** The *security set identifier (SSID)* is a name assigned to your wireless network. You should change the name of the SSID and also not use your company name. When hackers are war driving, they pick up on a signal from a wireless network, and if the SSID says “BridgetsWidgets,” they then look for the building with the “Bridget’s Widgets” sign. When

they spot the sign, they then drive close to the building so that they get a stronger signal. Don't make it easy for them to figure out what building to get close to!

- ◆ **Disable SSID broadcasting:** After you have set the SSID, you also want to disable the broadcasting of the SSID. The wireless router broadcasts the SSID so anyone who gets close will know the wireless network is there. If you disable broadcasting, then to connect to the wireless network, a person has to know and input the SSID manually into his or her network client.
- ◆ **Enable WEP:** *Wired Equivalent Privacy (WEP)* is a method to encrypt traffic from the wireless client to the wireless access point. If you enable WEP, you type a passphrase that is used to perform the encryption, and anyone who wishes to connect to your wireless network must also know this passphrase.

Getting an A+

This chapter introduces you to a number of best practices for securing your Windows environment. Some of the key points to remember for the exam are

- ◆ Create user accounts for each user of the system. Make sure that users use strong passwords for those accounts and understand not to share those passwords.
- ◆ Assign permissions to resources such as folders and files to ensure that unauthorized users don't get access to the resource.
- ◆ Enable auditing so that you are aware of any security-related events that happen on the system. Also be sure to review the security log often.
- ◆ Be sure to enable a firewall for the network and enable the firewall on the Windows system.
- ◆ Secure your wireless router by disabling the wireless feature if you are not using the wireless components of the router. If you are using wireless then be sure to implement WEP or WPA and disable SSID broadcasting.

Prep Test

- 1** What security feature stops network packets from entering the system through the network card?
 - A Auditing
 - B Password policy
 - C Permissions
 - D Firewall

- 2** What is the network name assigned to the wireless network?
 - A WEP
 - B SSID
 - C SID
 - D WPE

- 3** What permission on a folder is assigned to allow a user to read, modify, create, and delete a file?
 - A Read
 - B Full Control
 - C Modify
 - D Deny

- 4** You have enabled auditing; where do you go to view the audit information?
 - A Local Security Policy
 - B Event Viewer
 - C LAN Connection Properties
 - D Firewall

- 5** A privilege to perform an operating system task is known as what?
 - A Permission
 - B Policy
 - C Right
 - D Firewall

- 6** What permission allows a user to modify the permissions?
 - A Read
 - B Full Control
 - C Modify
 - D Deny

7 How would you allow Bob to change the time on his computer?

- A Enable an Audit Policy
- B Place Bob in the Administrators group
- C Assign Bob the change system time permission
- D Assign Bob the change system time right

8 Which security features might you enable through the system BIOS?

- A Boot devices
- B Password policy
- C Permissions
- D Audit policy

9 Which built-in group has full access to the system?

- A Administrator
- B Power Users
- C Account Operators
- D Administrators

10 Where do you go to enable the firewall in Windows?

- A Properties of the LAN connection
- B Local Security Policy
- C Event Viewer
- D Security page tab

Answers

- 1** D. A firewall is designed to stop data from entering your system through the network card. *See "Implementing Firewalls."*
- 2** B. The SSID is the name assigned to the wireless network. *Review "Securing wireless."*
- 3** C. To allow users to read, modify, create, and delete a file, you assign the Modify permission. *Check out "Implementing Permissions and Rights."*
- 4** B. When auditing has been enabled, you view the auditing information by reviewing the security log in Event Viewer. *Peruse "Implementing Auditing."*
- 5** C. A right gives you the privilege to perform an operating system task. *Take a look at "Rights."*
- 6** B. The Full Control permission allows users to modify permissions on a file or folder. *Peek at "Permissions."*
- 7** D. You would assign Bob the change system time right. You could put Bob in the Administrators group, but that is not the best answer because you have given him a number of other capabilities at the same time. *Look over "Rights."*
- 8** A. The boot devices can be disabled through the BIOS, which controls whether someone can bypass your operating system by booting from a bootable CD or floppy. *Study "Securing Systems through BIOS."*
- 9** D. Administrators is a group that is built-in that has full access to the system. *Refer to "Creating groups."*
- 10** A. The firewall can be enabled through the LAN connection properties. *Examine "Implementing Firewalls."*

Chapter 3: Viruses and Malicious Software

Exam Objectives

- ✓ Understanding viruses
- ✓ Looking at malicious software
- ✓ Performing software updates

In this chapter, you find out about the different types of malicious software that you might encounter after you connect your system to the Internet or start sharing files with another system. As a service technician, you will spend most of your time trying to fix a computer that is polluted with viruses or other forms of malicious software.

Being able to protect your system from viruses and other forms of malicious software is an important skill to have. You will also have to know how to perform a virus scan and remove viruses from your system.

My Computer Has a Virus!

A *virus* is a piece of software that is designed to do harm to the system in one fashion or another. The virus is typically associated with a file, and when that file is opened, the virus is activated and does its damage to the system. The damage could be disastrous in the sense that the system no longer starts up, or it may just create a faulty system that is annoying to work with. Either way, the virus has done its job — to cause pain and suffering in your computer's life, and yours!

Types of viruses

A number of different types of viruses can infect your system:

- ◆ **Boot sector virus:** A boot sector virus attacks the boot code contained in the boot sector of the disk. This type of virus loads every time the system boots up and sometimes may even prevent the system from booting.
- ◆ **File virus:** A file virus attaches itself to an executable file and is run, or activated, when the executable file is run. This has been one of the most

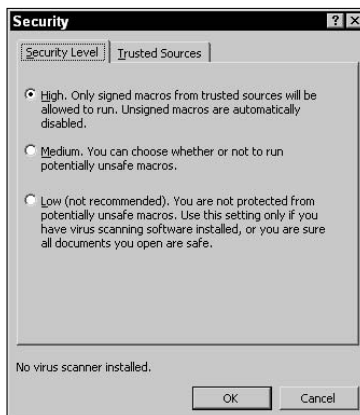
popular forms of viruses for many years because the attacker simply has to trick you into running the executable.

- ◆ **Macro virus:** A very popular type of virus today is the macro virus. A few years ago, Microsoft was promoting the fact that you could customize what Microsoft Office applications, such as Word, Excel, and Access, could do by programming your own macros. A macro is a set of Visual Basic for Applications (VBA) commands that performs a specific task. VBA is pretty much a full programming language.

The fact that VBA is a programming language is important because it can perform a lot of powerful actions, such as creating, deleting, and modifying files. VBA can also be used to call other programs; for example, I created a macro in Microsoft Excel that automatically starts Microsoft Outlook and e-mails the contents of the spreadsheet to one of my business partners.

After the power of VBA became known, hackers started creating malicious macros in Microsoft Word and Excel documents so that when you opened the documents, the macro would run and cause damage to your system. To prevent this from occurring, you could change the *macro security settings* in each of the Microsoft Office applications to control what you would like to do about macros found in documents when the document is opened. To change the macro security settings, choose Tools⇨Macro⇨Security in your Microsoft Office program and then choose one of the following options (shown in Figure 3-1):

Figure 3-1:
Configuring
macro
security
settings in
Microsoft
Office.



- **High:** Setting the macro security to high means that only macros from trusted sources are allowed to run, and all other macros are disabled. This is the highest level of security because most macro code is not created from trusted sources.

- *Medium:* Setting the macro security to medium means that the Microsoft Office application prompts you for whether you would like to enable or disable macro support every time you open the file. This is a good setting because if you were unaware that a file had a macro, you could choose to disable macros while the file is open — knowing that if there is malicious code in the file, it won't run.
- *Low:* Setting the macro security to low allows all macros to run without your knowledge. If you have no worries about macro security, you may set this option — but never open a Microsoft Office document unless you know where it came from!

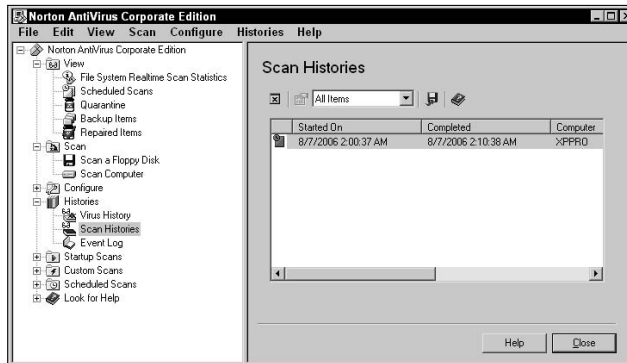


You may also encounter worm viruses or Trojan viruses, but I save those for discussion in “Understanding Malicious Software,” later in this chapter.

Protection from viruses

In order to protect your computers and servers from viruses, it is essential that you have virus-protection software installed. *Virus-protection software*, such as the program shown in Figure 3-2, is software that knows about the different viruses that exist and can either remove the virus or remove files from your system that contain viruses.

Figure 3-2:
Looking at
virus-
protection
software.



Just as many auto manufacturers provide different vehicles — and each vehicle has its own benefits — many different name brands offer virus-protection software. Some of the popular names in virus protection are

- ◆ McAfee Antivirus
- ◆ Norton Antivirus
- ◆ Panda Antivirus
- ◆ FProt Antivirus

Each antivirus software product has its own benefits and features. For example, you may like the interface or usability of one product over another. But at the same time, each of the products is very competitive and should have similar features available. When shopping for antivirus software, you should look for software that offers at least the following features:

- ◆ **Scheduled virus scans:** Virus scans can be scheduled to perform automatically on the system. This is a great feature because you can have the virus-scanning software scan the system in the middle of the night, when the system is not being used. You may also choose what happens when a virus is found — attempt to remove the virus from the file, place the file in a quarantine area, or delete the file. The benefit of a scheduled scan is that you do not have to physically perform the scan yourself.
- ◆ **Real-time protection:** Real-time protection is the idea that the virus-protection software runs in memory all the time and scans any file that you open. The benefit of real-time protection is that you are protected from viruses between the scan times.
- ◆ **Scheduled definition updates:** Virus definitions are what the virus-protection software uses to maintain its knowledge of what viruses are out there. The virus-protection software should have a feature that allows the definitions to be downloaded from the antivirus vendor's site.
- ◆ **Scanning e-mail:** Many versions of virus-protection software today support scanning e-mail messages as they arrive in your inbox. This is typical of virus-protection software that runs on an e-mail server.

Using virus-protection software

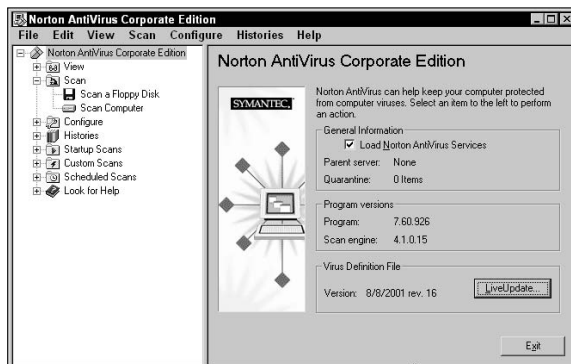
How you use virus-protection software differs from product to product, but the basic concepts should be the same. In this section, you find out how to navigate around Norton AntiVirus and how to execute tasks such as performing a virus scan and scheduling a scan to occur in the middle of the night.

Performing a virus scan

After you have your virus-protection software installed, you can perform a manual scan of your system. How you perform a manual virus scan, which detects and removes viruses from your system, will be different with each different product. I have Norton AntiVirus Corporate Edition, so I demonstrate the steps with it.

You start your antivirus software from the Start menu. For example, to start my Norton AntiVirus, I choose Start⇒All Programs⇒Norton AntiVirus Corporate Edition and then choose the Norton AntiVirus Corporate Edition program item. After you start your antivirus software, you see a screen similar to the one shown in Figure 3-3.

Figure 3-3:
Norton
AntiVirus
Corporate
Edition's
main
screen.



Notice in the main screen of the virus-protection software that I can perform a scan of a floppy disk or a scan of the computer by clicking the `scan computer` command on the left side of the screen. After you choose the `scan computer` command, you select which drives you wish to scan in the details pane and then click the Scan button to begin the scan.

Scheduling virus scans

Most virus-protection software lets you schedule virus scans for the middle of the night so that you don't have to waste time in the day performing the virus scan. To schedule a virus scan to occur in the middle of the night, start your virus-scanning software. In my example, I load up Norton AntiVirus Corporate Edition. After I have started Norton up, I expand Scheduled Scans on the left side and choose New Scheduled Scan. After you choose to create a new scheduled scan, you need to give the scheduled scan a name and description, as shown in Figure 3-4, and then choose Next. You are then asked when you would like to schedule the scan. I choose 2:00 a.m., every day, and then click Next. You are finally asked what drives to scan — select all drives and then click Save. You have successfully created a saved schedule so that virus scans on your system run every night at 2 a.m.



You can schedule a virus scan to occur at any time during the day. Overnight, however, is usually the best time to run the scan because it allows the scan to run without interfering with your use of the computer.

Enabling real-time protection

Virus-protection software today supports real-time protection, which prevents a virus from entering your system. Without real-time protection, you are susceptible to virus infection during the period of time between your manual or scheduled virus scans. With real-time protection, if you open a file that has a virus, the virus-protection software detects it and notifies you. To

enable real-time protection, start your virus-protection software and then expand the Configure option on the left side. You can then select the File System Real-time Protection option and ensure that real-time protection is enabled, as shown in Figure 3-5.

Figure 3-4: Setting the name and description for a scheduled scan.

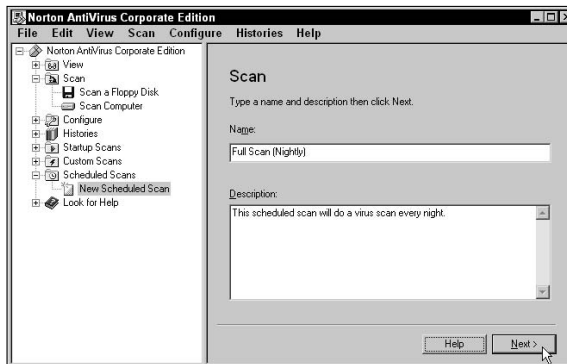
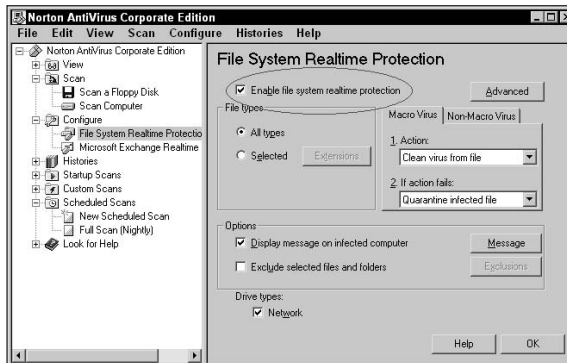


Figure 3-5: Enabling real-time protection within antivirus software.



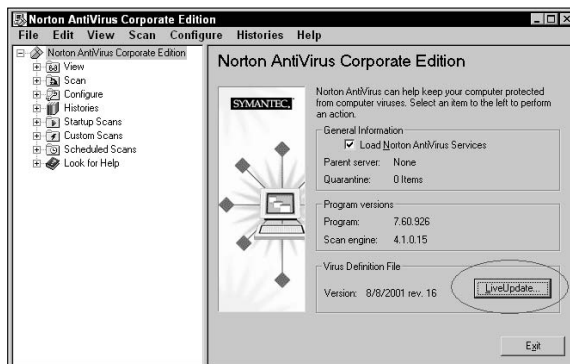
Updating virus definitions

When you perform a virus scan, your virus-protection software knows only about the viruses up to the creation time of the software. This is a huge problem because new viruses appear every day.

So, to keep your software valid and to allow it to still be useful years after you purchase it, the manufacturer of the virus-protection software uses *virus definitions* as a way for the software to know the current list of viruses. The virus definitions can be updated through the virus-protection software. If you update the virus definitions, it means that although your software may be two years old, it is current as far as the viruses it can protect you against.

It is critical that you update your virus definitions regularly, or your virus-protection software won't know about any new viruses that are developed. To update the virus definitions in my Norton AntiVirus, after I start the software, I see a LiveUpdate button on the main screen, as shown in Figure 3-6. To update my virus definitions, I simply click the LiveUpdate button; updating virus definitions should be this simple in most other antivirus programs as well. It is also important to notice that the antivirus software also displays how up-to-date the virus definitions are. In my example, I have not run LiveUpdate yet, so my software only knows about viruses developed before August 2001. I really need to run the update!

Figure 3-6:
Updating
virus
definitions
with the
LiveUpdate
feature.



Understanding Malicious Software

Other types of malicious software, outside of your typical viruses, attack systems every day. The following sections outline other types of malicious software — but understand that they are all considered types of viruses.

Trojan horses

Trojan horses are special programs that do something totally different than what the user who runs it thinks it does. For example, NetBus is a very popular Trojan virus that ships as a file called `patch.exe`. A hacker e-mails the file called `patch.exe` and explains in the e-mail that this is a security patch you need to apply to make sure your system is secure. Unfortunately, `patch.exe` is the security hole! When you run `patch.exe`, it opens your system up to the hacker by opening a port so that the hacker can connect to the port at any time and control your system.

Trojan viruses are normally loaded on your system by the hacker tricking you into running the program on the system. You can remove the Trojan with virus-protection software.

Worms

A *worm* is a self-replicating virus. By *self-replicating*, I mean that the worm doesn't need to be activated by the user opening the file. A worm is a virus that runs on a system and also tries to infect other systems on the network. The Nimda virus is an example of a worm virus.

Worms are loaded on your system by connecting to your system from across the Internet. The worm is usually designed to infect the system by connecting through a specific piece of software. For example, Nimda was designed to infect any system running Microsoft Web server software called *Internet Information Services (IIS)*. To prevent a worm from infecting your system, be sure to have a firewall and make sure that you are up to date with patches. For more information on firewalls check out Book IX, Chapter 2, and to learn more about staying up-to-date read the "Preventing problems by staying up to date" section in this chapter. To remove a worm from your system, you can try your antivirus software, or you may have to download a removal tool specific for that worm. For example, Microsoft offers a malicious code removal tool that you can download from www.microsoft.com/downloads.

Spyware and adware

Spyware is software that loads on your system and then monitors your Internet activity, while adware is software that creates pop-ups from time to time advertising a particular product or service.

Both of these types of viruses infect your system when you surf the wrong Internet site. Spyware and adware have become a huge negative result of the Internet, so a number of products are available to eliminate spyware and adware. The most popular products used to eliminate spyware and adware are

- ◆ **Spybot Search & Destroy**
- ◆ **Ad-Aware**
- ◆ **Microsoft's Windows Defender**

When selecting which software to use for spyware or adware removal, be sure to go with a product that supports features similar to good antivirus software, such as real-time protection and scheduled spyware scans. Microsoft's Windows Defender (shown in Figure 3-7) is a free download at <http://www.Microsoft.com/downloads> that supports these features along with advanced features that protect your browser from being hijacked.

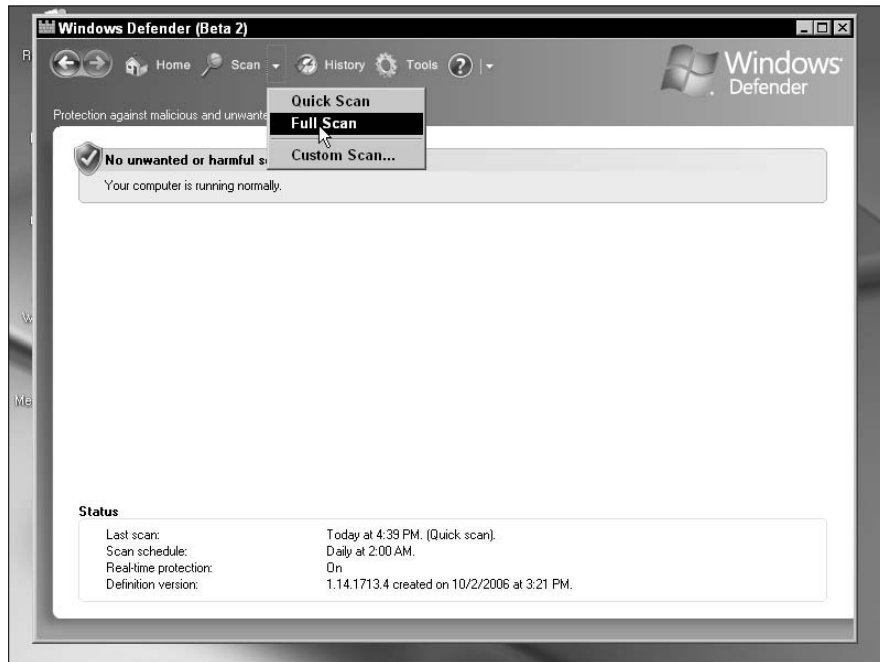


Figure 3-7:
Removing
spyware
with
Microsoft's
Windows
Defender.

Identifying Hoaxes

It is important to understand that some malicious code that you hear about is a *hoax*! For example, I remember a few years back, I was playing the elf bowling program that everyone was e-mailing around at Christmas time. After weeks of playing this game, it was said that the program should be removed from your system because on a certain date it would do damage to your system. I removed the program from my main computer, but I ran the program on a test system after that date, and it did no harm.

You typically receive hoaxes about viruses through your e-mail system. The e-mail you receive is acting as a virus alert, but unfortunately, there is no actual virus to report — it is a hoax. The benefit of such hoaxes is for the creators of the *actual* viruses; hackers hope that you receive so many hoaxes that you eventually ignore true virus alerts.

If you receive an e-mail or other form of notice about a virus, you should check it against a virus hoax list to see whether the warning message is a hoax or if it has merit. Most virus-protection software manufacturers keep an up-to-date list. For example, McAfee lists hoaxes at <http://vil.mcafee.com/hoax.asp>. You may also go to a generic hoax site like Hoaxbusters, which is found at <http://hoaxbusters.ciac.org>.

Preventing Problems by Staying Up to Date

One of the most popular techniques hackers use to compromise systems is to find vulnerabilities in the software that we use day in and day out. For example, hackers quickly figured out a way to perform the “dot dot” attack on Windows 2000 systems after IIS was installed. With the “dot dot” attack, hackers navigate the folder structure of a Web server and delete files — a serious security flaw. After Microsoft got wind of the mistake, its programmers created a fix. It is your responsibility as a network administrator to download all the fixes to problems in the software you use.

Windows Update

To make it easier for you to get *security fixes*, also known as *patches*, and updates to Microsoft software, Microsoft created the Windows Update feature within the operating system. If you choose the Windows Update command from the Start menu, you are automatically connected to the Microsoft Windows Update site, where your system is scanned for which updates are needed.

Windows Update allows you to do an express update where all critical updates are installed on your system, or a custom update (see Figure 3-8) where you get to select which updates to install (see Figure 3-8). From the Windows Update site, you install all the updates or patches that your system needs. Performing a Windows Update is a critical step to securing your systems — be sure to do it regularly.



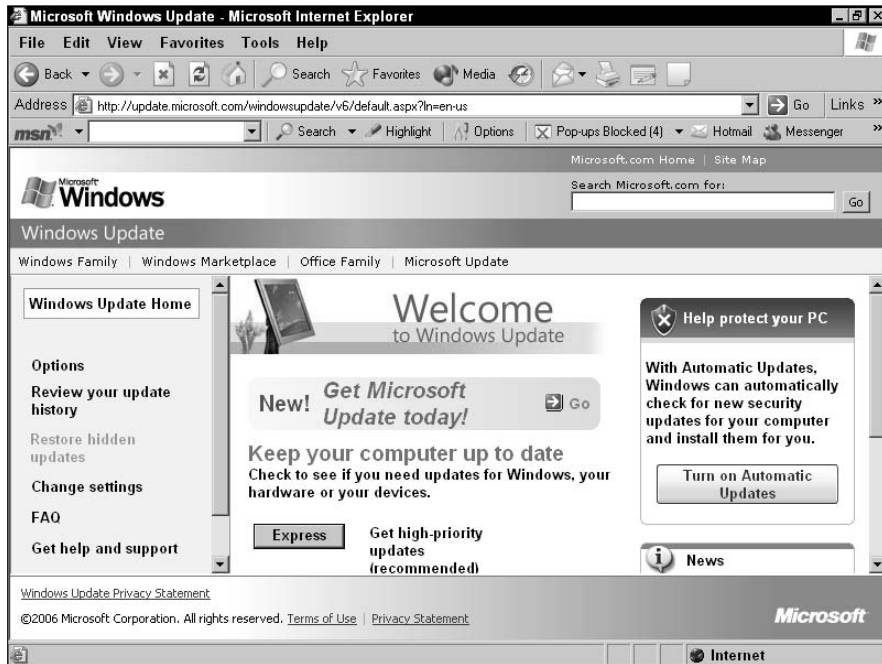
Microsoft has changed from Windows Update to Microsoft Update so that you can now get updates for more than just the Windows operating system. You can download updates for a number of Microsoft products from the Microsoft Update site, such as Windows and Microsoft Office.

SUS and WSUS

Most companies have thousands of systems on the network, which would make it impossible to visit each system individually and run Windows Update. To solve the problem of deploying updates to thousands of systems at a time, Microsoft created *Software Update Service (SUS)*. The SUS software is loaded on a server and gets all of the updates from the Windows Update site. After the SUS server downloads all the updates, you review and approve the updates. The SUS software then sends the updates to the other systems on your network. This means that you don't have to run Windows Update on every system on your network — saving you time and money!

The software update service was improved a few years later to create *Windows Software Update Services (WSUS)*. WSUS solves a number of problems that existed in the original SUS. Some of the benefits of WSUS over SUS are

Figure 3-8:
The
Windows
Update site
allows you
to do an
express
update or
a custom
update.



- ◆ **Synchronization:** *Synchronizing* is downloading updates from the Microsoft Update site to the SUS server. When you synchronize the SUS server with Windows Update, it takes hours because all of the updates that exist on the Windows Update site are downloaded to your SUS server. This makes no sense because you receive updates for software you don't use — wasting time and space on your system. WSUS downloads a listing of updates when you choose to synchronize the WSUS server. After you get the listing of updates, you then select which updates to download — saving space and time!
- ◆ **Product selection:** With WSUS, you can select which Microsoft products you use so that you get a listing of updates for only those products. For example, if you aren't using Microsoft Exchange server, it doesn't make sense to get updates for that product. With WSUS, you can specify that you are not using that software, and the updates for that software do not appear in the list.
- ◆ **Computer Groups:** With WSUS, you can deploy updates to selected groups of systems. This is a big improvement over SUS because with SUS, when you approve the update, the update gets sent to all the systems. With WSUS, you can create groups of computers, and when you approve an update, you assign it to a group. For example, you could create a group called "Accounting" and then place all the accounting

computers in that grouping, which allows you to approve updates for just the accounting systems if you like.

Antivirus

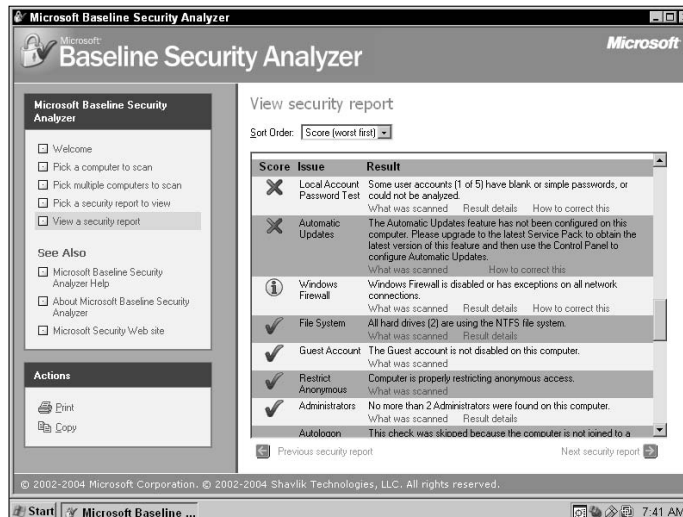
A big part of keeping your system up to date is making sure that you download security fixes from the Windows Update site or from the WSUS server, but you also need to be sure that you update the virus definitions for your virus-protection software. You want to make sure that you investigate how to update the virus definitions with your antivirus software. The virus-protection software should also give you the opportunity to schedule virus definition updates.

Other security software

Outside of virus-protection software and Windows Update, you should be aware of a number of other software products.

One of the most beneficial types of software out there for network administrators is vulnerability assessment software. *Vulnerability assessment* software can scan all the systems on your network and let you know if you are missing any security patches or if you are breaking any major security best practices, like not having a password on your administrator account. Microsoft has a vulnerability scanner called *Microsoft Baseline Security Analyzer (MBSA)* (shown in Figure 3-9), which you can download from www.microsoft.com/downloads.

Figure 3-9:
MBSA identifies any security issues with your system.



The MBSA identifies a number of issues surrounding the security of your system. A few examples of the types of information you will be presented with are

- ◆ **Check for Windows Vulnerabilities:** MBSA verifies that your current system is up to date with operating system patches. It also identifies any security best practices that have been broken.
- ◆ **Weak Passwords:** MBSA ensures that all the user accounts have strong passwords.
- ◆ **IIS and SQL Server Vulnerabilities:** MBSA verifies that if you are using IIS and SQL servers that you are up to date with security patches specific to those products.

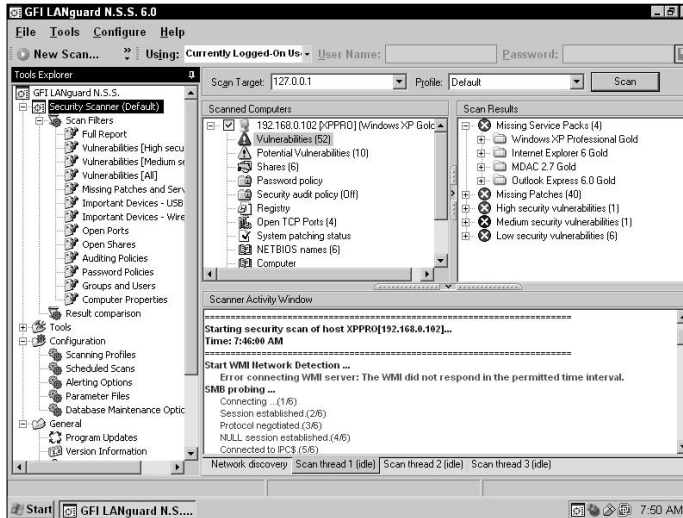
Another example of a vulnerability scanner that I use quite often is GFI's *LANguard Network Security Scanner*. LANguard is very similar to MBSA but offers a lot more detail with regards to the overall picture of the network and how secure the systems are. LANguard (shown in Figure 3-10) has the benefit of logging the information to a Microsoft Access database or SQL server so that you can create your own reports on the information collected. LANguard collects vulnerability information but also reports the following information:

- ◆ **Information on the password policy**
- ◆ **Groups that exist**
- ◆ **User accounts that exist**
- ◆ **Services that are running**
- ◆ **Shares that exist**
- ◆ **NetBIOS name information**



To practice performing a Microsoft Update to your system, check out Lab 3-1, and to practice using the Microsoft Baseline Security Analyzer (MBSA), check out Lab 3-2. Lab 3-1 and Lab 3-2 can be found in the `Labs.pdf` file in the `Author` directory of the CD-ROM.

Figure 3-10:
GFI's
LANguard
identifies
security
issues with
the
operating
system.



Getting an A+

This chapter introduces the concepts of viruses and malicious software. Some key points to remember when you take the A+ Certification exam are

- ◆ A *virus* is a program that does harm to your system.
- ◆ Popular types of viruses are file viruses, boot sector viruses, and macro viruses.
- ◆ Other types of malicious software are Trojan horse viruses, worms, and spyware.
- ◆ You may use antivirus software to remove viruses from your system and anti-spyware software to remove spyware from your system.
- ◆ Be sure to update your virus definitions on a regular basis so that your virus-protection software can protect your system against the most popular viruses.

Prep Test

- 1** What must be updated in your antivirus software?
 - A Virus software
 - B Virus definitions
 - C Patches
 - D Trojans

- 2** What type of malicious software does something totally different than what the user expects?
 - A Virus
 - B Virus definitions
 - C Worm
 - D Trojan virus

- 3** What type of virus can prevent the system from booting?
 - A File virus
 - B Macro virus
 - C Boot sector virus
 - D Trojans

- 4** What type of virus might occur in a Microsoft Office document?
 - A File virus
 - B Macro virus
 - C Boot sector virus
 - D Trojans

- 5** What type of malicious software monitors your Internet activity?
 - A Adware
 - B Worm
 - C Spyware
 - D Trojans

- 6** What type of malicious software creates pop-ups on your system?
 - A Adware
 - B Worm
 - C Spyware
 - D Trojans

Viruses and Malicious Software

- 7** What is the best way to ensure that your system is up to date with security patches?
- A Virus-protection software
 - B Virus definition update
 - C Spyware
 - D Windows Update
- 8** What software allows the network administrator to deploy patches to all the systems on the network after they are downloaded to a central server?
- A Virus-protection software
 - B Virus definitions
 - C WSUS
 - D Trojans

Answers

- 1** B. Virus definitions are needed to ensure that your virus-protection software knows about recent viruses. *See "Updating virus definitions."*
- 2** D. A Trojan virus is a piece of software that does something totally different than what the user expected when he installed the application. *Review "Trojans."*
- 3** C. A boot sector virus prevents a system from booting because it corrupts the boot sector. *Check out "Types of viruses."*
- 4** B. A macro virus is created in a programming language like VBA, which exists in the Microsoft Office suite. *Peruse "Types of viruses."*
- 5** C. Spyware is the type of malicious software that monitors Internet activity. *Take a look at "Spyware and adware."*
- 6** A. Adware is the type of malicious software that is responsible for creating pop-ups on your system. *Peek at "Spyware and adware."*
- 7** D. You should perform a Windows Update regularly to download current patches to your system. *Look over "Windows Update."*
- 8** C. WSUS can deploy updates to all systems on the network. *Study "SUS and WSUS."*

Appendix A: About the CD

On the CD-ROM:

- ✓ **Quick Assessment tests to help you gauge your strong and weak areas of knowledge**
- ✓ **Dummies Test Engine with hundreds of sample questions to make sure you're ready for the A+ Certification exam**
- ✓ **Lab Manual with exercises to help you practice and hone your skills**

System Requirements

Make sure that your computer meets the minimum system requirements shown in the following list. If your computer doesn't meet most of these requirements, you may have problems using the software and files on the CD. For the latest and greatest information, please refer to the ReadMe file located at the root of the CD-ROM.

- ◆ A PC with a Pentium or faster processor; or a Mac OSX computer with a G3 or faster processor
- ◆ Microsoft Windows 2000 or later; or Mac OSX system software 10.1 or later
- ◆ A CD-ROM drive

If you need more information on the basics, check out these books published by Wiley Publishing, Inc.: *PCs For Dummies*, by Dan Gookin; *Macs For Dummies*, by David Pogue; *Windows XP For Dummies* and *Windows 2000 Professional For Dummies*, both by Andy Rathbone.

Using the CD

To install the items from the CD to your hard drive, follow these steps.

- 1. Insert the CD into your computer's CD-ROM drive. The license agreement appears.**

Note to Windows users: The interface won't launch if you have autorun disabled. In that case, click Start→Run. In the dialog box that appears,

type **D:\start.exe**. (Replace D with the proper letter if your CD-ROM drive uses a different letter. If you don't know the letter, see how your CD-ROM drive is listed under My Computer.) Click OK.

Note for Mac Users: The CD icon will appear on your desktop, double-click the icon to open the CD and double-click the "Start" icon.

- 2. Read through the license agreement, and then click the Accept button if you want to use the CD.**
- 3. The CD interface appears. The interface allows you to install the programs and run the demos with just a click of a button (or two).**

What You'll Find on the CD

The following sections are a summary of the software you'll find on the CD-ROM included with this book.

Quick Assessments

The Quick Assessment tests are designed to help you get comfortable with the A+ testing situation and pinpoint your strengths and weaknesses on the topic. These tests are located within the Quick Assessments folder in the Author directory. There are tests for the individual chapters of each book, and taking these tests before reading the chapter will help you determine what areas you should spend extra time on when studying for the exams.

To utilize this tool, click on the document for the chapter you are about to read. You can choose to answer the questions as you read them on-screen, or you can print the document and answer on paper. After you take the Quick Assessment test, you will know which answers you got right and which ones you got wrong — immediately after the questions is a section with the answers. Included with the answer to each question is a cross-reference to the specific section in that chapter where you can find the text that answers the question.

Dummies Test Engine

The Dummies Test Engine is designed to simulate the actual A+ situation — a question with multiple choice answers. The Dummies Test Engine on the CD-ROM is not adaptive and it is not timed — but you may time yourself to gauge your speed. After you answer each question, you will find out whether or not you answered the question correctly. If you answered correctly, you are on your way to A+ success. If you answered incorrectly, you will be told the correct answer with a brief explanation of why it is the correct answer.

The Dummies Test Engine includes all of the Prep Test questions from the end of each chapter in the book, as well as hundreds of additional questions. If you perform well on the Dummies Test Engine, then you're probably ready to tackle the real thing.

Lab Manual

The Lab Manual included on the CD-ROM is designed to give you hands-on experience executing specific A+-related tasks. Doing these labs will help you become comfortable with the tasks you will be performing on-the-job, and they will also allow you to hone your skills. These labs are located within the Author directory in the `Labs.pdf` file. The labs in this document are organized by Book and Chapter so that you can easily locate them.

As you read through the chapters within this book, you will come across text marked with an "On The CD" icon. To access the labs referenced by the text, click on the `Labs.pdf` file and then locate the Book, Chapter, and lab you are looking for using the file's Table of Contents. You can choose to execute the labs as you read them on-screen, or you can print the entire `Labs.pdf` file — which you may want to do as many of the labs require you to record information on a piece of paper.

Troubleshooting

I tried my best to compile programs that work on most computers with the minimum system requirements. Alas, your computer may differ, and some programs may not work properly for some reason.

The two likeliest problems are that you don't have enough memory (RAM) for the programs you want to use, or you have other programs running that are affecting installation or running of a program. If you get an error message such as `Not enough memory` or `Setup cannot continue`, try one or more of the following suggestions and then try using the software again:

- ◆ **Turn off any antivirus software running on your computer.** Installation programs sometimes mimic virus activity and may make your computer incorrectly believe that it's being infected by a virus.
- ◆ **Close all running programs.** The more programs you have running, the less memory is available to other programs. Installation programs typically update files and programs; so if you keep other programs running, installation may not work properly.
- ◆ **Have your local computer store add more RAM to your computer.** This is, admittedly, a drastic and somewhat expensive step. However, adding more memory can really help the speed of your computer and allow more programs to run at the same time.

If you have trouble with the CD-ROM, please call the Wiley Product Technical Support phone number at (800) 762-2974. Outside the United States, call 1(317) 572-3994. You can also contact Wiley Product Technical Support at <http://support.wiley.com>. John Wiley & Sons will provide technical support only for installation and other general quality control items. For technical support on the applications themselves, consult the program's vendor or author.

To place additional orders or to request information about other Wiley products, please call (877) 762-2974.

Appendix B: CompTIA A+ Exam Reference Matrix

You can use this chart to identify which chapters to study in preparation for the CompTIA A+ Certification Exams. Because the topics for each exam have a number of overlapping areas, this matrix can be used to point you to correct chapters for each exam. Since you need to take two exams, then it is still best to read this entire volume because there is a great deal of overlap between the Essentials exam and all of the other exams.

After reading the entire volume, you can use this chart as part of your review. If you go through the objectives and find any area you don't think you know, you can immediately turn to the appropriate chapter for in-depth review.

The exams covered in the matrix and their column key identifiers are:

- ◆ **220-601 Essentials (E)**
- ◆ **220-602 IT Technician (I)**
- ◆ **220-603 Remote Support Technician (R)**
- ◆ **220-604 Depot Technician (D)**

2006 Examination Objectives

<i>Domain/Objective/ Description</i>	<i>E</i>	<i>R</i>	<i>I</i>	<i>D</i>	<i>Book-Chapter</i>
1.0 Personal Computer Components	✓	✓	✓	✓	
1.1 Identify the fundamental principles of using personal computers	✓				
Identify the names, purposes, and characteristics of storage devices	✓				2-5
Identify the names, purposes, and characteristics of motherboards	✓				2-1
Identify the names, purposes, and characteristics of power supplies, such as AC adapter, ATX, proprietary, voltage	✓				2-6
Identify the names, purposes, and characteristics of processor/CPUs	✓				2-2

1088 2006 Examination Objectives

<i>Domain/Objective/ Description</i>	<i>E</i>	<i>R</i>	<i>I</i>	<i>D</i>	<i>Book-Chapter</i>
Identify the names, purposes, and characteristics of memory	✓				2-3
Identify the names, purposes, and characteristics of display devices, such as projectors, CRT and LCD	✓				3-3
Identify the names, purposes, and characteristics of input devices, such as mouse, keyboard, bar code reader, multimedia (e.g., Web and digital cameras, MIDI, microphones), biometric devices, touch screen	✓				3-2, 3-6
Identify the names, purposes, and characteristics of adapter cards	✓				3-2, 3-3, 3-6
Identify the names, purposes, and characteristics of ports and cables, for example, USB 1.1 and 2.0, parallel, serial, IEEE-1394/FireWire, RJ-45 and RJ-11, PS/2/Mini-DIN 6, Centronics (e.g., mini, 36), multimedia (e.g., 1/8" connector, MIDI, COAX, S/PDIF)	✓				3-1
Identify the names, purposes, and characteristics of cooling systems, such as heat sinks, CPU and case fans, liquid cooling systems, thermal compound	✓				2-2
1.2 Install, configure, optimize and upgrade personal computer components	✓	✓	✓	✓	
Add, remove, and configure internal and external storage devices	✓			✓	2-5
Install display devices	✓				3-3
Add, remove, and configure basic input and multimedia devices	✓				3-3, 3-6
Add, remove, and configure display devices, input devices, and adapter cards, including basic input and multimedia devices		✓			3-3, 3-6
Add, remove, and configure personal computer components, including selection and installation of appropriate component			✓	✓	2-1, 2-2, 2-3, 2-4, 2-5, 2-6, 3-2, 3-3, 3-4, 6-1
1.3 Identify tools, diagnostic procedures, and troubleshooting techniques for personal computer components	✓	✓	✓	✓	

<i>Domain/Objective/ Description</i>	<i>E</i>	<i>R</i>	<i>I</i>	<i>D</i>	<i>Book-Chapter</i>
Recognize the basic aspects of troubleshooting theory	✓				1-2, 4-2
Recognize names, purposes, characteristics and appropriate application of tools			✓	✓	4-2
Identify and apply basic diagnostic procedures and troubleshooting techniques	✓	✓	✓	✓	1-2, 4-2
Recognize and isolate issues with display, power, basic input devices, storage, memory, thermal, POST errors (e.g., BIOS, hardware)	✓				4-2
Recognize and isolate issues with display, peripheral, multimedia, specialty input devices, and storage		✓			4-2
Recognize and isolate issues with peripherals, multimedia, specialty input devices, internal and external storage, and CPUs			✓		4-2
Identify steps used for and apply basic troubleshooting techniques to check for problems (e.g., thermal issues, error codes, power, connections including cables and/or pins, compatibility, functionality, software/drivers, proper seating, installation, appropriate component, settings, current driver) with components	✓		✓	✓	4-2
Apply steps in troubleshooting techniques to identify problems (e.g., physical environment, functionality, and software/driver settings) with components including display and input devices and adapter cards		✓			4-2
Recognize the names, purposes, characteristics and appropriate applications of tools, for example, BIOS self-test, hard drive self-test, and software diagnostic test	✓				4-2
1.4 Perform preventative maintenance on personal computer components	✓	✓	✓	✓	
Identify and apply basic aspects of preventative maintenance theory	✓	✓			4-1

1090 2006 Examination Objectives

<i>Domain/Objective/ Description</i>	<i>E</i>	<i>R</i>	<i>I</i>	<i>D</i>	<i>Book-Chapter</i>
Identify and apply common preventative maintenance techniques for devices	✓		✓	✓	4-1
2.0 Laptops & Portable Devices	✓		✓	✓	
2.1 Identify the fundamental principles of using laptops and portable devices	✓		✓	✓	
Identify names, purposes, and characteristics of laptop-specific components and accessories	✓			✓	3-7
Identify and distinguish between mobile and desktop motherboards and processors, including throttling, power management, and WiFi	✓				3-7
Identify appropriate applications for laptop-specific communication connections, such as Bluetooth, infrared, cellular WAN, and Ethernet			✓		3-7
Identify appropriate laptop-specific power and electrical input devices and determine how amperage and voltage can affect performance			✓	✓	3-7
Identify the major components of the LCD, including inverter, screen, and video card			✓	✓	3-7
2.2 Install, configure, optimize, and upgrade laptops and portable devices	✓		✓	✓	
Configure power management	✓				3-7
Demonstrate safe removal of laptop-specific hardware, such as peripherals, hot-swappable devices, and non-hot-swappable devices	✓		✓	✓	3-7
Describe how video sharing affects memory upgrades			✓	✓	3-7
2.3 Identify tools, diagnostic procedures, and troubleshooting techniques for laptops and portable devices	✓		✓	✓	
Use procedures and techniques to diagnose power conditions, video issues, keyboard and pointer issues and wireless card issues	✓		✓	✓	4-2
2.4 Perform preventative maintenance on laptops and portable devices	✓				

<i>Domain/Objective/ Description</i>	<i>E</i>	<i>R</i>	<i>I</i>	<i>D</i>	<i>Book-Chapter</i>
Identify and apply common preventative maintenance techniques for laptops and portable devices, for example, cooling devices, hardware and video cleaning materials, operating environments (including temperature and air quality), storage, transportation, and shipping	✓				4-1
3.0 Operating Systems — unless otherwise noted, operating systems referred to include Microsoft Windows 2000, XP Professional, XP Home, and Media Center only	✓	✓	✓		
3.1 Identify the fundamentals of using operating systems	✓	✓	✓		
Identify differences between operating systems (e.g., Mac, Windows, Linux) and describe operating system revision levels, including GUI, system requirements, application and hardware compatibility	✓				5-1
Identify names, purposes, and characteristics of the primary operating system components, including registry, virtual memory, and file system	✓				5-2
Describe features of operating system interfaces	✓				5-2
Identify the names, locations, purposes, and characteristics of Windows 2000, XP Professional, and XP Home system files	✓				5-6
Identify concepts and procedures for creating, viewing, and managing disks, directories, and files in operating systems	✓	✓	✓		5-4
Use command-line functions and utilities to manage Windows 2000, XP Professional, and XP Home, including proper syntax and switches		✓	✓		5-5, 8-3
3.2 Install, configure, optimize, and upgrade operating systems — references to upgrading from Windows 95 and NT may be made	✓	✓	✓		
Identify procedures for installing operating system	✓				5-3

1092 2006 Examination Objectives

<i>Domain/Objective/ Description</i>	<i>E</i>	<i>R</i>	<i>I</i>	<i>D</i>	<i>Book-Chapter</i>
Identify procedures for upgrading operating systems	✓				5-3
Install/add a device, including loading, adding device drivers, and installing required software	✓				6-1
Identify procedures and utilities used to optimize operating systems	✓	✓	✓		6-2, 6-3
Locate and use operating system utilities and available switches		✓	✓		4-1, 5-2, 5-5, 6-4, 7-3
3.3 Identify tools, diagnostic procedures and troubleshooting techniques for operating systems	✓	✓	✓		
Identify basic boot sequences, methods, and utilities for recovering, and the ability to recover operating system	✓		✓		7-3
Identify and apply diagnostic procedures and troubleshooting techniques	✓				1-2
Recognize and resolve common operational problems	✓	✓	✓		6-2, 7-3
Recognize, explain, and resolve common error messages and codes	✓	✓	✓		7-2
Identify the names, locations, purposes and characteristics of operating system utilities	✓	✓			4-1, 5-5, 6-4, 7-3
Use diagnostic utilities and tools to resolve operational problems		✓	✓		1-2, 5-2, 6-4, 7-2, 7-3
3.4 Perform preventative maintenance on operating systems	✓	✓	✓		
Describe and demonstrate common utilities for performing preventative maintenance on operating system, for example, software and Windows updates (e.g., Service Packs), scheduled backups/restore, restore points	✓	✓	✓		7-3
4.0 Printers and Scanners	✓	✓	✓	✓	
4.1 Identify the fundamental principles of using printers and scanners	✓	✓	✓	✓	
Identify differences between types of printer and scanner technologies (e.g., laser, inkjet, thermal, solid ink, impact)	✓	✓			3-5, 3-6

<i>Domain/Objective/ Description</i>	<i>E</i>	<i>R</i>	<i>I</i>	<i>D</i>	<i>Book-Chapter</i>
Describe processes used by printers and scanners, including laser, ink dispersion, thermal, solid ink, and impact printers and scanners			✓	✓	3-5, 3-6
Identify names, purposes and characteristics of printer and scanner components (e.g., memory, driver, firmware) and consumables (e.g., toner, ink cartridge, paper)	✓				3-5, 3-6
Identify the names, purposes, and characteristics of interfaces used by printers and scanners, including port and cable types	✓				3-1, 3-5, 3-6, 8-2
4.2 Installing, configuring, optimizing, and upgrading printers and scanners	✓	✓	✓	✓	
Identify the steps used in the installation and install and configure printers/scanners	✓	✓	✓	✓	3-5, 3-6
Install and configure printer/scanner upgrades, including memory and firmware			✓	✓	3-5, 3-6
Optimize printer performance, for example, printer settings such as tray switching, print spool settings, device calibration, media types, and paper orientation	✓				3-5
Optimize scanner performance, for example, resolution, file format, and default settings		✓	✓		3-6
4.3 Identify tools, diagnostic procedures, and troubleshooting techniques for printers and scanners	✓	✓	✓	✓	
Gather information about printer/scanner problems	✓	✓	✓	✓	3-5, 3-6
Review and analyze collected data	✓		✓	✓	1-2
Isolate, identify solutions, and resolve printer/scanner problems	✓		✓	✓	1-2
Troubleshoot a print failure (e.g., lack of paper, clear queue, restart print spooler, recycle power on printer, inspect for jams)		✓			3-5
Identify appropriate tools for troubleshooting and repairing printer/scanner problems			✓	✓	3-5, 3-6, 4-2

<i>Domain/Objective/ Description</i>	<i>E</i>	<i>R</i>	<i>I</i>	<i>D</i>	<i>Book-Chapter</i>
4.4 Perform preventative maintenance of printers and scanners			✓	✓	
Perform scheduled maintenance according to vendor guidelines (e.g., install maintenance kits, reset page counts)			✓	✓	3-5, 3-6
Ensure a suitable environment			✓	✓	4-2
Use recommended supplies			✓	✓	3-5, 3-6
5.0 Networks	✓	✓	✓	✓	
5.1 Identify the fundamental principles of networks	✓	✓	✓		
Describe basic networking concepts	✓				8-1
Identify names, purposes, and characteristics of network cables (e.g., RJ-45 and RJ-11, ST/SC/LC fiber connectors, USB, IEEE-1394/FireWire)	✓				3-1, 8-1
Identify names, purposes, and characteristics (e.g., definition, speed, and connections) of technologies for establishing connectivity	✓	✓	✓		8-1, 8-2, 8-4
Identify names, purposes, and characteristics of the basic network protocols and terminologies		✓	✓		8-3, 8-4
5.2 Install, configure, optimize, and upgrade networks	✓	✓	✓		
Install and configure network cards (physical address)	✓				8-3
Install, identify, and obtain a wired and a wireless connection	✓				8-2, 8-3
Establish network connectivity		✓	✓		8-3
Demonstrate the ability to share network resources		✓	✓		8-1, 8-3
Install and configure browsers			✓		8-4
5.3 Identify tools, diagnostic procedures, and troubleshooting techniques for networks	✓	✓	✓		
Explain status indicators, for example, speed, connection and activity lights, and wireless signal strength	✓				8-2, 8-3

<i>Domain/Objective/ Description</i>	<i>E</i>	<i>R</i>	<i>I</i>	<i>D</i>	<i>Book-Chapter</i>
Identify the names, purposes, and characteristics of command-line tools		✓	✓		8-1, 8-3
Diagnose and troubleshoot basic network issues		✓	✓		8-1, 8-3, 8-4, 9-2
5.4 Perform preventative maintenance of networks, including securing and protecting network cabling			✓		9-2
6.0 Security	✓	✓	✓	✓	
6.1 Identify the fundamental principles of security	✓	✓	✓	✓	
Identify names, purpose, and characteristics of hardware and software security	✓				9-1
Identify the names, purposes, and characteristics of access control and permissions		✓	✓		9-2
Identify names, purpose, and characteristics of wireless security	✓				8-2, 9-2
Identify names, purpose, and characteristics of data and physical security	✓			✓	9-2
Describe importance and process of incident reporting	✓				9-2
Recognize and respond appropriately to social engineering situations	✓				9-1
Identify the purposes and characteristics of auditing and event logging			✓		9-2
6.2 Install, configure, upgrade, and optimize security	✓	✓	✓	✓	
Install, configure, upgrade, and optimize hardware, software, wireless, and data security components	✓	✓	✓	✓	9-2
6.3 Identify tools, diagnostic procedures, and troubleshooting techniques for security	✓	✓	✓		
Diagnose and troubleshoot hardware, software and data security issues	✓	✓	✓		9-2
6.4 Perform preventative maintenance for computer security	✓	✓	✓		

1096 2006 Examination Objectives

<i>Domain/Objective/ Description</i>	<i>E</i>	<i>R</i>	<i>I</i>	<i>D</i>	<i>Book-Chapter</i>
Implement software security preventative maintenance techniques, such as installing Service Packs and patches and training users about malicious software prevention technologies	✓				9-2, 9-3
Recognize social engineering and address social engineering situations		✓	✓		9-1
7.0 Safety and Environmental Issues	✓		✓	✓	
7.1 Describe the aspects and importance of safety and environmental issues	✓				
Identify potential safety hazards and take preventative action	✓				1-3
Use Material Safety Data Sheets (MSDS) or equivalent documentation and appropriate equipment documentation	✓				1-3
Use appropriate repair tools	✓				1-3
Describe methods to handle environmental and human (e.g., electrical, chemical, physical) accidents including incident reporting	✓				1-3
7.2 Identify potential hazards and implement proper safety procedures, including ESD precautions and procedures, safe work environment, and equipment handling	✓				1-3
7.3 Identify proper disposal procedures for batteries, display devices, and chemical solvents and cans	✓				1-3
7.4 Identify potential hazards and proper safety procedures, including power supply, display devices, and environment (e.g., trip, liquid, situational, atmospheric hazards and high-voltage and moving equipment)			✓	✓	1-3
8.0 Professionalism and Communication	✓	✓	✓		
8.1 Use good communication skills, including listening and tact/discretion, when communicating with customers and colleagues	✓	✓	✓		1-2

<i>Domain/Objective/ Description</i>	<i>E</i>	<i>R</i>	<i>I</i>	<i>D</i>	<i>Book-Chapter</i>
8.2 Use job-related professional behavior, including privacy, confidentiality, and respect for the customer and customer's property	✓	✓	✓		
Behavior	✓		✓		1-2
Property	✓		✓		1-2

Index

Symbols and Numerics

- * wildcard, 615, 622
- ? wildcard, 616, 622
- / ? switch, 627–628, 658
- 2.4 GHz spectrum, 911, 912, 919
- 3DNow! Technology, 136, 137
- 8.3 naming convention, 580, 581
- 10 Mbps network card, 326
- 10/100 Mbps network cards, 326, 432
- 12-volt power connector, 276, 278
- 16-bit Windows applications, 544, 710–716, 822
- 24-pin Molex connector, 276, 277
- 32-bit processor, 116
- 32-bit protected-mode state, 517
- 32-bit registers, 118
- 32-bit Windows applications, 706–710
- 40-wire ribbon cable, 217, 218
- 64-bit processors, 134–135
- 64-bit version of Windows, 707
- 640KB memory in MS-DOS, 662
- 80387 chip, 120
- 80486 chips, 119

A

- A+ Certified Professional certification, 1
- A+ Depot Technician Exam, 12, 13
- A+ domains (exam areas), 11

- A+ Essentials Exam, 10, 11–12
- A+ exams, cost of taking, 14
- A+ IT Technician Exam, 12
- A+ Remote Support Technician Exam, 12–13
- absolute path, 617, 619
- AC (Alternating Current), 41
- AC (Alternating Current) adapters, 280–281, 422, 426–427
- accelerated video cards, 337
- access control, setting, 959
- access method of a URL, 998–999
- access points, wireless, 908, 910
- access time, 211, 514
- access token, 1024–1025
- accessibility features, 320–321
- account management, auditing for, 1052
- Account Operators group, 1043
- ACE (Advanced Computing Environment) initiative, 649
- ACPI (Advanced Configuration and Power Interface), 426
- Active Directory Database, 1024, 1052
- active heat sink, 142
- active matrix display, 428
- active partition, 241, 244, 761
- actuator arm, 208
- adapter cards, 53
- Adapter Unit Interface (AUI) port, 882
- adaptive exams, 10, 11, 16
- Add a Network Place Wizard, 526
- Add Hardware Wizard, 682

- Add or Remove Programs, 528, 569, 701–703
- Add Printer Wizard, 392, 969
- Add/Remove Hardware, 527, 686
- Add/Remove Windows Components, 569–570
- address bus, 117–118, 130
- address lines, 117, 525
- Address Resolution Protocol (ARP), 953–954
- address space, 709, 710
- addressing in SCSI, 225–227
- Ad-hoc mode for a wireless network, 909
- administrator account, 552, 1041, 1042
- administrator password assigning for a wireless router, 1059
- changing for a wireless network, 915
- setting in CMOS, 192, 1040
- in a Windows 2000 installation, 552
- in a Windows XP installation, 559
- Administrators group, 1043
- ADSL (Asymmetric Digital Subscriber Line), 58, 1002–1003
- Advanced Configuration and Power Interface (ACPI), 426
- Advanced Encryption Standard (AES), 917
- Advanced Graphics Port. *See* AGP
- Advanced Micro Devices (AMD), 136–138
- Advanced Packaging Tool (APT), 504–505
- Advanced Power Management (APM), 426

- Advanced Programmable Interrupt Controller (APIC), 358
- Advanced RISC Computing Specification. *See* ARC pathnames
- Advanced Technology Attachment (ATA) standard, 215
- adware, 1072
- AES (Advanced Encryption Standard), 917
- AGP (Accelerated Graphics Port)
 - bus, 57, 103, 104, 335
 - slot, 74, 75, 77, 94
- airflow, 450–454, 457
- alerts in System Monitor, 744–745
- allocation unit, 237, 252
- alphanumeric keys on a keyboard, 316
- alternate location, restoring to, 849–850
- Alternating Current. *See* AC
- AMD (Advanced Micro Devices), 136–138
- amperes (amps), 272
- AMR (Audio/Modem Riser) bus architecture, 105
- analog audio cable, 481
- analog audio connector, 310
- analog line, 323
- analog multimeters, 468
- analog video and cable connectors, 310
- Andreessen, Marc, 990
- animated logo, disabling, 656
- answer file, 561–562
- antistatic bags, 40
- anti-static mats, 469
- anti-static wrist strap, 469, 470
- antivirus software. *See* virus-protection software
- APIC (Advanced Programmable Interrupt Controller), 358
- APIPA (Automatic Private IP Addressing), 946
- APM (Advanced Power Management), 426
- application(s)
 - crashes, information about, 780–781
 - defined, 700
 - installing and removing, 699–704
 - multithreaded, 121, 706
 - not displaying complete filenames, 581
 - not installing, 823
 - not starting, 821–823
 - running 16-bit Windows, 710–716
 - running 32-bit Windows, 706–710
 - running in separate memory spaces, 714–715
 - standalone, 700
 - terminating from memory, 711
 - terminating under Windows XP, 714
 - testing on an operating system, 543–544
- Application log
 - in Event Viewer, 777, 778
 - of Windows XP, 806
- application-layer protocols, 985
- Applications tab of Task Manager, 714, 741
- APT (Advanced Packaging Tool), 504–505
- ARC pathnames, 649–650
- architecture of operating systems, 517–518
- Archive attribute, 595, 596, 606
- archive bit, 1029–1030
- arguments in batch files, 632–633
- ARP (Address Resolution Protocol), 953–954
- ARPANET, 981–982, 987
- ASR (Automatic System Recovery) process, 853–855
- Asynchronous Digital Subscriber Line (ADSL), 58, 1002–1003
- AT, IRQs available on, 357
- AT commands, 324, 325
- AT connector, 318. *See also* DIN keyboard connector
- AT keyboard, 317
- ATA (AT Attachment), 51
- ATA/ATAPI-6 specification, 480
- ATAPI specification, 215–216
- Athlon, 126, 127, 137
- Athlon XP, 126, 127, 137
- Athlon XP M, 136
- attachments with e-mail messages, 987
- attacks, types of, 1013–1020
- attended installations, 544–561
- attrib command, 625
- attrib.exe command, 603–606, 833
- ATX main power connector, 276–279
- ATX motherboard
 - ATX power connector on, 84
 - compared to MicroATX motherboard, 96
 - described, 90, 93–95
 - I/O ports on, 74, 75
 - P4 power connector on, 84

power supplies common to, 274

USB ports on, 80, 81

ATX power connectors, 83–84, 274

ATX power supply, 274

ATX v2.2 form factor, 279

audio cues, 475–476

audio problems, 482

Audio/Modem Riser (AMR), 105

auditing, 241, 1051–1054, 1058

AUI (Adapter Unit Interface) port, 882

authentication, 1023–1025

authorization, 1025–1026

AutoComplete settings, 997

autoexec.bat

calling, 654

described, 834

modifying TSR lines in, 668

to override default settings, 794

overriding default, 719–720

path variable setting, 637

processing, 659–660

automatic link tracking, 821

Automatic Private IP Addressing (APIPA), 946

Automatic System Recovery (ASR) process, 853–855

automatic update tools, 450

automatic updates, causing auto restarts, 817

AutoPlay feature of Windows, 703

auto-restart errors, 815–817

autorun.inf file, 703

AutoScan option of msdos.sys, 655

B

B channels in ISDN, 1003

Baby AT motherboard, 90, 91–92

back probing connectors, 489, 490

background printing, 394

Background settings for MS-DOS applications, 726

BackTrack, 921

backup, 844, 845–849, 846–849, 1030

backup and restore strategy, 1029

Backup Operators group, 1043

backward compatibility, 699, 709, 728

bad CD-ROM device, 564

bad clusters, 251

Bad command or file name message, 637

bad memory, 816

bad sectors, 558

bar code reader, 329

barrel connectors, 308, 309

Base 2 numbers, 580

Base 16 numbers, 819

baseband transmission, 886

bash (Bourne Again Shell), 993

basic disks, 244–248, 625

Basic Input-Output System. *See* BIOS

basic rate ISDN, 1003

batch files, 590, 631–638, 659

batteries

disposal of, 42–43

laptop, 422, 423–426

recommended care procedures, 461

troubleshooting, 487–489

battery for CMOS RAM, 73, 74, 155–156

baud rate, 292, 485

BEDO (Burst Extended Data Output), 158

Berners-Lee, Tim, 990, 991

bidirectional communication, 292

bidirectional port mode, 189

Big Drives specification, 480

binary number system, 580

biometric devices, 330

biometrics, 1023

BIOS

built-in testing software, 474

date and revision number, 194–195

described, 55, 152–153

on the motherboard, 73

potential issues with upgrading, 181

purpose of, 179–180

recognizing larger drives, 212

reserving resources, 692

securing systems, 1039–1040

supporting LBA or ECHS, 212

supporting Plug and Play, 690

supporting USB, 486

troubleshooting, 476

updating in EEPROM, 153

upgrading, 181–183

BIOS chip

on the motherboard, 72, 180

updating, 72, 182–183

BIOS password, 1040

bitmap fonts, 720

bits, 116, 118

- bits per second (bps),
292, 323, 485
- .bks (backup selections)
file, 847
- blackouts, 282, 456
- blank pages from laser
printers, 401
- blanking plate, 436
- blowhole, 453
- Blue Screen of Death, 517
- Blue-Screen errors, 817–819
- Bluetooth, 433, 919–920
- BNC connector, 303–304,
879–880
- BNC-T connector, 304, 881
- body of a network packet,
1054
- boot devices
 - CMOS control of, 1040
 - disabling, 1021
- boot disks, 38, 474, 538,
831–835
- boot errors, 793–795
- boot files, 796–799, 842, 846
- boot floppy for Windows,
537–538
- boot keys in Windows 9x,
655
- boot loader for Windows XP,
647
- boot loader section, 648
- Boot Partition, 647
- boot process, 183–184
- boot sector, 646
- boot sector virus, 1065
- boot sequence, CMOS entry
for, 191
- bootable CDs, 474, 538, 547
- bootable devices in SCSI,
230
- bootable disk devices, 646
- booting, 196, 645, 836–840
- boot.ini file, 647, 648,
796, 797–798
- bootlog, 802, 838–839

- bootlog.txt file,
837, 838–839
- bottlenecks, 735–736
- box cutter, 473
- bps (bits per second),
292, 323, 485
- bridge, 893–894
- broadband cable, 58
- broadcast messages, 953
- brownouts, 282
- Browser service, 972–973
- browsers, 990, 994–998
- browsing network
resources, 966–967
- brute force attack, 1016
- BSD OS, Mac OS X based
on, 507
- BSD UNIX, 505
- BSoD (Blue Screen of
Death) errors, 817–819
- BSSID (Basic Service Set
Identifier), 910
- build numbers, 516
- built-in accounts,
1041–1042
- built-in coprocessors, 337
- built-in groups in Windows,
1043
- built-in network adapter,
196
- burning a CD, 234
- Burst Extended Data
Output (BEDO), 158
- bus
 - architectures, 96–106
 - definition of, 51
 - described, 116
 - mastering, 98, 99, 101
 - providing video services,
56–57
 - topology, 871–872
- Bus Enumerators, 691, 692
- Bush, Vannevar, 989
- buttons, removing
rendering, 729

C

- .cab files, 700, 784, 785
- cable access to the
Internet, 1002
- cable adapters, 308–310
- cable orientation, 298–299
- cable problems with AC
adapters, 427
- cable television, cable for,
878
- cable testers, 473, 884
- cable types, 227, 294–297
- cables, 457, 473
- cabling, 227–228, 232,
876–884
- cache memory, 118–119
 - described, 69–70
 - disabling in CMOS,
197–198
 - on IDE devices, 214
 - levels of, 164
 - on motherboards, 106
 - SRAM used for, 154–155
 - working with, 164–166
- cache space, 738
- caching high-speed RAM, 52
- Cacti tool, 460
- calibration routine, for
printers, 463
- call home functionality,
disabling, 803
- call to a memory address,
117
- call-waiting, 484
- camera resolution, 411–412
- cameras, digital and web,
411–416
- capacitive keyboard, 316
- capacitor input power
supplies, 283
- capacitors, 41
- card services, 435
- cartridges, disposal of, 43

- cascading controllers, 357–358
- cascading DMA channel, 362
- cascading hubs, 896
- case(s)
 - elements inside, 50–56
 - form factors for, 56
 - maintenance of, 457
 - moisture-proofing, 454
 - taking apart, 472
- case standoffs, 471
- CAT 3 UTP cabling, 886
- CAT 5 cabling, 297, 877, 884
- CAT 5e cabling, 877
- CAT 5 UTP cabling, 297, 887, 888
- CAT 6 cabling, 297, 877
- Cathode Ray Tube monitors. *See* CRT monitors
- CCD (charge-coupled device), 408, 409, 411
- CD(s)
 - adding to the EIDE chain, 214
 - BIOS update program
 - from, 181
 - booting and running
 - the operating system
 - from, 38
 - booting from, 191, 538, 547, 1040
 - damaged by glue in labels, 461
 - existing on an ATA chain, 216
 - sharing in Windows XP, 965
 - speed (transfer rate) of, 234
 - writing to, 235
- CD drivers, locating, 835
- CD drives, 234, 460, 480–481
- CD Recordable (CDR), 234
- CD ReWriteable (CDRW), 234–235
- CD writers on an ATA chain, 216
- CDR (CD Recordable), 234
- CD-ROM (for this book), 2, 14
- CDRW (CD ReWriteable), 234–235
- Celeron, 127, 131–132
- cell phones, EMI caused by, 455
- cellular, networking or Internet access, 919
- Central Processing Unit. *See* CPU
- centralized administration, 869–870
- Centronics connectors, 77, 226, 302
- .CER file, 601
- CERN (European Council for Nuclear Research), 983
- certificates
 - digital, 1028
 - public keys stored in, 599
 - SSL, 992
 - SSL protocol, 997
 - used by EFS recovery, 601
 - used for security in *ssh*, 994
- certification, 10
- CFX v1.3 (Compact Form Factor), 279, 280
- CGA standard, 336
- Change permission, 963–964
- channels
 - in T1, 1003
 - for wireless networks, 912
- character repeat options, 320
- charge-coupled device (CCD), 408, 409, 411
- chdir* (*cd*) command, 619
- Check Disk utility, 261–262
- /checkupgradeonly* switch, 566, 567
- chip creep, 457
- chip extractor, 36
- chip inserter, 37
- chip packages, 122–125
- chipsets, 54
- chkdsk* in Recovery Console, 843
- chkdsk.exe*, 450, 480, 624, 764–766, 833
- choice.exe*, 635–636
- cipher.exe* command-line utility, 600–601
- classes of IP addresses, 941, 942
- clean memory location, 666
- cleaners, 450
- cleaning blade, 382, 384
- cleaning cloths, 449
- cleaning phase of laser printing, 384
- Clear To Send (CTS), 307
- cleartext, 993, 1027
- client(s)
 - in a peer-to-peer network, 867
 - reducing per access point, 920
 - in wireless networks, 908
- Client for Microsoft Networks, 932–934
- Client Service for NetWare, 932–935
- client-server networking, 869–871
- clock cycles, 116, 727
- clock multiplying, 128
- clusters
 - bad, 251
 - described, 210
 - listed in the FAT, 251
 - size decreased in FAT32, 239–240
 - used inefficiently by the FAT file system, 237

- `cmd.exe`, 613–615
- CMOS
 - backing up before a password is set, 193
 - changing settings in, 1039
 - compared to BIOS, 179–180
 - configuring parallel port mode in, 189
 - configuring startup sequence in, 191
 - described, 55–56, 72–73, 155, 183–185
 - parallel or serial ports disabled in, 485
 - reserving resources in, 370
 - security settings in setup program, 1040
 - System Summary screen, 188
 - troubleshooting, 476
 - virus protection setting in, 197
- CMOS battery, 184, 185, 190, 192, 193, 424
- CMOS configuration, 23, 155–156
- CMOS password, 1021, 1022
- CMOS RAM, 73, 155–156, 184–185
- CMOS Save & Restore utility, 193
- CMOS settings, 185–193, 194, 222, 458, 544
- CMOS setup program, 179, 180, 183, 185, 186
- CNR bus architecture, 106
- coaxial (coax) cabling, 878–882
- CODECs, 415–416
- cold swapping, 692
- collision detection, 885
- Color Graphics Adapter.
See CGA standard
- color inkjet printers, 385
- color LCDs, 428
- colors
 - incorrect on monitors, 483
 - supported by video cards, 337–338
- COM ports, 75, 432–433
- COM+ Class Registration database, 846
- COM1, 190, 354, 360
- COM2, 190, 354, 360
- COM3, 354, 360
- COM4, 354, 360
- combo card, 328
- command files, 590
- command interpreter, 658
- command line, 527, 603–606
- command line tools, 623
- Command Prompt boot, 840
- command prompt, help
 - from, 658
- Command Prompt window, 603
- command shells, 993–994
- `command.com`, 658–659
 - bad or missing, 794–795
 - described, 613–615, 833, 834
 - loading low, 656
 - passing control to, 654
- command-line interface, added to Mac with the BSD-like kernel, 508
- command-line utilities, 613–615
- commands, locating, 637
- commit charge in Task Manager, 738
- communication
 - components of laptop computers, 432–433
 - methods of network devices, 898
 - skills, 27–28
 - in `telnet`, 993
- Communication and Network Riser (CNR), 106
- communications (COM) ports. *See* COM ports
- Compact Disc-Read Only Memory. *See* CD(s)
- `compact.exe`, 598–599
- companion CD, 2, 14
- comparisons, making, 622–623
- Complementary Metal-Oxide Semiconductor. *See* CMOS
- components
 - adding or removing from Windows, 568–570
 - disposing of PC, 42–43
 - maintenance of, 457–458
 - replacing opposed to repairing, 475
 - testing related, 24–25
- Compress attribute, 597–599
- compressed air, 449, 457
- compressed file, 599, 785
- compression, 241, 252, 591
- Compression/Decompression software, 415–416
- CompTIA, 9, 14
- CompTIA A+ Certification, 9–13
- computer(s)
 - backing up, 845–849
 - basic functions of, 49–50
 - identifying changes to, 21
 - keeping on or off, 456
 - navigating, 523–527
 - networking, 898–900
 - restoring to a previous state, 852–853
 - restoring to default factory installation, 856
 - reverting to a previous state, 850

- computer browser, 972
- computer chip, 38–39
- Computer Information
 - section in Network Diagnostics, 780
- Computer Management
 - utilities or snap-ins in, 775
 - in Windows 2000 or Windows XP, 624, 774–776
- Computer Management console
 - creating a user account in, 1040–1041
 - creating groups, 1044–1045
 - creating user accounts, 1042–1043
- Computer Management window, 244
- Computer Name Changes dialog box, 947–948
- computer names. *See also* NetBIOS names
 - entering in a Windows 2000 installation, 552
 - planning for the Windows OS, 541
 - resolving to IP addresses, 954
- computer ports. *See* ports
- computer settings in the Registry, 523
- computer vacuum, 449
- condensation, 451, 454
- conditioning phase of laser printing, 383
- confidentiality, respecting, 29
- config.sys
 - described, 656–657, 834
 - error in a line of, 794
 - modifying, 664
 - MS-DOS environment settings from, 654
- overriding default, 719–720
 - seeing all lines in, 794
 - used by Windows 9x, 794
- configuration change, testing, 23
- Configuration Manager, 691, 692
- configuration utilities, 687, 766–776
- connectors
 - checking, 476
 - for external SCSI devices, 226
 - for memory modules, 167
 - on the motherboard, 78–79, 82–88
 - for multimedia, 310
 - types of, 299–306
- context setting in Client Service for Netware, 934
- continuity test in a multimeter, 469
- continuity tester, 488
- continuous form feed printers, 380
- Control Panel, 527–529
- controller(s)
 - for hardware, 70
 - integrated into drives, 214
- controller bus as SCSI or multi, 796
- conventional memory
 - described, 662–663
 - making available to applications, 722
 - for MS-DOS applications, 721–723
 - protecting, 723
- cookies, 996–997
- cooling fans, 54, 142, 143, 273
- cooperative multitasking, 705, 710
- coprocessors, 337
- copy backups, 848
- copy command, 620–621
- core operating system files, 846
- corrosion on the motherboard, 454
- corrupted memory space, 822
- counter logs in System Monitor, 745–746
- counters in System Monitor, 742, 743
- CPU (Central Processing Unit), 50–51, 66, 115. *See also* processors
- CPU fan, 815–816
- CPU usage, 736, 737
- CPU utilization problems, 740
- crash dump files, 780, 803, 804
- create volume command, 247
- credentials, 847
- critical resources, 735–736, 743–744
- cross-linked files, 765
- crosstalk, 387
- CRT (Cathode Ray Tube) monitors, 41–42, 43, 340. *See also* monitors
- cscript.exe, 590, 784
- CSMA/CA, 885, 912
- CSMA/CD, 884–885
- CTS (Clear To Send), 307
- current, 273, 468
- cursor blink rate, 320
- cursor keys on a keyboard, 316
- customers, 27–28
- “cut-off” corner of a socket, 140, 141
- cylinders, 210
- CYMK cartridges, 385

D

- D channel in ISDN, 1003
- DAC (digital-to-analog) converter, 336, 343
- daemons, 992
- daily backups, 848
- daisy chain, 225, 230, 409
- daisy chaining
 - defined, 189
 - by EPP, 189, 346
 - by FireWire, 82
 - by SCSI, 346
 - by USB devices, 80
- damaged PIF, 822
- DAT (Digital Audio Tape), 237
- data
 - accessing in memory, 157
 - backing up, 1029–1031
 - destroying, 1029
 - importing into the Registry, 800
 - loss of, 447–448
 - protection of, 1028–1032
 - securing with RAID, 252–259
- data bus, 116
- Data Carrier Detected (DCD), 307
- data integrity, 260
- data lanes, used by PCI Express, 105
- data path of SIMM modules, 160
- Data Set Ready (DSR), 307
- data skew, 387, 388
- Data Terminal Ready (DTR), 307
- database of problems and resolutions, 25
- dataflow analysis, 130
- date and time in CMOS, 190, 194–195
- date and time, verifying, 553
- daughter board, 53
- DB9-male serial port, 76
- DB25-female parallel port, 76
- DB25-male serial port, 76
- DB-9 connector, 300, 301
- DB-15 connector, 294, 300, 301
- DB-25 connector, 300–301, 302
- DB-25 female connector, 228
- DB-25 male connector, 228
- DC (Direct Current), 41
- DC adapters, 280
- DC controller, 422
- DC voltage regulation, 273
- DCD (Data Carrier Detected), 307
- DDR (Double Data Rate) memory, 159
- DDS (Digital Data Storage), 237
- debug.exe, 833
- Debugging Mode option, 840
- default gateway, 943, 956
- default IRQ (interrupt request), 344
- default user template, 773
- defrag utility, 623–624
- defrag.exe, 450, 583, 763
- defragmentation, 762–764
- defragmentation utility, 260–261
- del command, 620, 621–622
- delete command in Windows setup, 542
- Demilitarized Zones (DMZs), 917, 1056–1057
- demodulation, 58, 323
- denial of service (DoS) attacks, 1016–1017
- Dependencies tab, 754
- deploy.cab file, 562
- desktop, launching, 526
- desktop metaphor, 499, 502
- DET (Directory Entry Table), 250, 251
- Details view in My Computer, 525
- developing phase of laser printing, 383
- device(s)
 - accessing memory directly, 360–362
 - assigned a unique IRQ line number, 356
 - configured with I/O addresses, 352–353
 - defined, 352
 - files referencing, 800
 - prioritizing by IRQ value, 359
 - supported by IDE and EIDE, 214
 - supported by SCSI chain, 223
 - troubleshooting installation, 695
 - during the troubleshooting process, 844
 - types supported by SCSI chain, 223
 - viewing hidden, 800
- device conflicts, 365–371
- device drivers
 - deleting installed, 687
 - described, 681
 - failure of, 694–696
 - installing, 682–684
 - loading and removing, 681–686
 - permissions to install, 693–694
 - removing, 682
 - resolving errors, 821

- signed and unsigned, 689–690
- updating, 688–689
- upgrading, 799
- verifying installation, 694
- device load process, 651–652
- Device Manager
 - accessing resource settings, 685
 - changing device settings, 687
 - changing resource assignments, 369
 - checking for missing drivers, 564
 - in Computer Management, 775
 - configuring network cards, 328
 - described, 774
 - locating and deleting devices, 800
 - removing devices, 686
 - reserving resources with, 370–371
 - serial or parallel ports disabled in, 485
 - updating a network driver, 929–930
 - verifying driver installation, 694
 - viewing DMA channels, 362
 - viewing I/O addresses, 353–355
 - viewing memory addresses, 363
 - viewing system resources, 364
- device related errors, 799–801
- device settings, changing, 687–688
- dfrg.msc, 763
- DHCP (Dynamic Host Configuration Protocol), 918, 945–946, 971
- Diagnosis tab of Dr. Watson, 803
- diagnostic software, 473–474
- diagnostic tools, 474, 803–806
- diagnostic utilities, 776–782
- diagnostics at the command prompt, 623–624
- diagonal cutter, 473
- dialog boxes, MS-DOS-based, 628
- dial-up access to the Internet, 1002
- dictionary attack, 1015–1016
- differential backups, 848, 1030
- digital audio connection, 310
- Digital Audio Tape (DAT), 237
- digital cameras, 411–412
- digital certificate, 1028. *See also* certificates
- Digital Data Storage (DDS), 237
- digital infrared thermometers, 471
- digital multimeters, 37, 468–469
- digital signal, converting to analog, 58
- Digital Signal Processor (DSP), 322
- Digital Versatile Disks (DVDs), 235
- digital video cameras, 414–415
- digital zoom, 412
- digital-to-analog converter (DAC), 336, 343
- digitizer, 492
- dimensions for a hard drive, 185–186
- DIMM sockets, 69
- DIMMs (Dual Inline Memory Modules), 161–162, 163, 167
- DIN keyboard connector, 79, 91, 294
- DIN-5 connector, 317, 318
- DIP (Dual Inline Package) chips, 70, 123
- DIP switch, 89
- dir command, 615–618
- dircmd environment variable, 618
- Direct Current (DC), 41
- direct hardware calls, 716
- Direct Memory Access channels. *See* DMA channels
- directories, 615–622
- Directory Entry Table (DET), 250, 251
- Directory Services Restore Mode boot option, 840
- Direct-Sequence Spread Spectrum technology, 912
- Disable Automatic Restart on System Failure option, 840
- disable command in Recovery Console, 844
- Disable Visual Themes setting, 729
- Disable Volume Shadow Copy option, 847
- disasters, preparing for, 831–861
- Disk Cleanup, 262–263, 750
- disk contention, 480
- Disk Defragmenter in Computer Management, 260, 775

- disk drives. *See also* hard drives
 - adding to printers, 391
 - defragmenting, 747
 - maintenance of, 259–263
 - physically destroying, 1029
 - types of, 244–249
- disk duplexing, 253, 254
- disk errors, 480
- disk geometry, 207–210
- disk maintenance, utilities for, 759–766
- Disk Management
 - creating primary partitions, 761
 - creating remaining partitions, 543
 - described, 625–626
 - preparing disks to hold data, 759–762
- Disk Management snap-in in Computer Management, 775
 - described, 244
 - formatting, 251
- disk mirroring, 253, 254
- disk partitioning, 625, 649, 760
- disk performance, 211
- disk platters, 208–209, 210
- disk quotas in NTFS 5.0, 240–241
- disk size limits of BIOS versions, 479–480
- disk stripping, 253–256
- disk-caching program, 669
- diskcomp.com, 622
- diskcopy command, 622
- diskpart, 843–844
- diskpart.exe, 625–626
- display settings
 - changing for applications, 729
 - configuring, 341–342
- displays, types of, 340–341
- distracters in exam questions, 15
- distractions, avoiding, 30
- distributed administration model, 869
- distribution, creating a Linux, 504
- distribution point, 571–572
- DivX CODEC, 415
- DLL (Dynamic Link Library) files, 781
- DMA (Dynamic Memory Access) channels, 292, 360–362
- DMA mode, 217
- DMZs (Demilitarized Zones), 917, 1056–1057
- DNS (Domain Name System), 951–952, 958, 1000–1001
- DNS servers, 944, 951, 1000
- docking stations, 433
- Document Type Definitions (DTDs), 991
- documentation, 26–27, 368
- domain, 541–542, 553, 554–555
- domain controller, 1024, 1052
- domain names, 1000
- Domain Name System. *See* DNS
- DoS (denial of service) attacks, 1016–1017
- DOS Mode, 840
- DOS programs, testing, 544
- DOS protected-mode interface (DPMI) memory, 724
- doskey.com, 659
- “dot dot” attack, on Windows 2000 systems, 1074
- dot matrix printers, 380, 386–387, 400
- dot pitch, 340
- dots per inch (DPI), 749
- dotted decimal notation format, 940
- Double Data Rate (DDR) memory, 159
- double density disks, 478
- Double Space driver (`dblspce.bin`), 656
- Dr. Watson, 780–781, 803, 804
- draft-quality printers, 387
- DRAM (Dynamic RAM), 154, 156–159
- drive connectors on motherboards, 85–88
- drive faceplates, 473
- drive imaging, 539–540
- drive overlay programs, 479
- Drive Space driver (`drvspace.bin`), 656, 833, 834
- drive type in CMOS, 186
- Drivelmage Pro, 539
- drivers
 - changing for printers, 395
 - loaded in conventional memory, 722
 - loading the network cards, 899
 - signing, 689–690
 - software, 928
 - types of, 85
 - updates, 450
 - upgrading for printers, 748
- drives. *See* hard drives
- `drstsn32.exe`, 780
- drum scanner, 408
- `drvspace.bin`, 656, 833, 834
- `drwtsn32.exe`, 803
- D-shell connectors, 300, 306
- DSL, 1003
- DSP (Digital Signal Processor), 322
- DSR (Data Set Ready), 307

- DSSS (Direct-Sequence Spread Spectrum), 912
- DTDs (Document Type Definitions), 991
- DTR (Data Terminal Ready), 307
- dual booting, 570
- dual core processors, 121, 122
- Dual Inline Memory Modules. *See* DIMMs
- Dual Inline Package (DIP) chips, 70, 123
- dual inline package (DIP) switch, 89
- dual-boot system, 238–239
- dual-ported memory, 156
- dual-ported VRAM, 338
- dual-scan display, 428
- Duron processor, 126, 127, 137
- DV format, 415
- DVD drives, 216, 480–481
- DVDs (Digital Versatile Disks), 235
- DVI (Digital Visual Interface) port, 294
- dynamic disks, 247–249, 761
- dynamic execution, 130
- Dynamic Host Configuration Protocol. *See* DHCP
- Dynamic Memory Allocation, 725
- Dynamic RAM (DRAM), 154, 156–159
-
- E**
-
- eavesdropping attack, 1017–1018
- ebd.cab, 834
- ebd.sys, 833
- ECC (error checking and correction) memory, 164
- ECHS (Extended Cylinder/Head/Sector), 212–213
- ECP (Extended Capabilities Port), 189, 292, 346
- edit command, 628–630
- Edit File Type dialog box, 587, 588
- edit.com
- changing config.sys, 794
- described, 784, 833
- opening .reg files in, 772
- .pifs created for it, 718
- EDO (extended data output) memory, 158, 166
- EEPROM (Electrically Erasable Programmable ROM), 72, 153
- EEPROM (Electrically Erasable Programmable ROM) chip, 180, 181, 368
- EFS (Encrypting File System), 240–241, 600, 601
- EGA standard, 336
- EIDE standard
- compared to IDE, 214
- connectors, 92
- described, 215, 216
- devices, 214
- drives, 221
- EIDEIATA-2 standard, 215
- EISA (Extended Industry Standard Architecture), 98–99
- elective exams, 10
- electrical interference, 884
- electrical tape, 473
- Electrically Erasable Programmable ROM. *See* EEPROM
- ElectroMagnetic Interference (EMI), 40, 455–456
- ElectroStatic Discharge. *See* ESD
- e-mail, 987–989, 1068
- Emergency Repair Disk (ERD), 769, 841
- Emergency Repair Process, 841–843
- EMI (ElectroMagnetic Interference), 40, 455–456
- emm386.exe, 663, 667–669
- employers, value of the A+ Certification to, 10
- emulator for expanded memory, 663
- enable command in Recovery Console, 844
- Enable command to re-enable a network connection, 931
- encoding, 987
- Encrypt attribute, 597, 599–601
- encrypted files, 600
- encrypting, 1027, 1028
- Encrypting File System, 240–241, 600, 601
- encryption, 600, 916
- end of file (EOF) marker, 251
- end-user licensing agreement (EULA), 545
- Energy Star specification, 273
- Enhanced 101-key keyboard, 429
- Enhanced Graphics Adapter (EGA), 336
- Enhanced Integrated Drive Electronics. *See* EIDE
- enhanced keyboard, 317
- Enhanced Parallel Port (EPP), 189, 292, 346
- environment variables, 618, 634, 659, 722

- environmental controls, 450–457
 - EOF (end of file) marker, 251
 - EPA (U.S. Environmental Protection Agency), 42
 - EPP (Enhanced Parallel Port), 189, 292, 346
 - EPROM (Erasable Programmable ROM), 52, 153
 - equality, testing, 634
 - equipment, high-voltage, 40–42
 - erase command, 621–622
 - ERD (Emergency Repair Disk), 769, 841
 - ergonomic keyboard, 317–318
 - ergonomics, 44–45
 - error checking memory, 163–164
 - error correction in Ultra DMA devices, 216
 - error logs in Dr. Watson, 781
 - Error Reporting tool, 803–805
 - error-checking and correction (ECC) memory, 164
 - errorlevel, testing, 635
 - errors, diagnostic tools for, 803–806
 - ESD (ElectroStatic Discharge)
 - avoiding, 38–40, 469
 - defined, 39
 - described, 454–455
 - ESSID (Extended Service Set Identifier), 910
 - Ethernet, 432, 886–888, 890
 - even parity, 163, 164, 485
 - event log, saving, 779
 - Event Log service, 805
 - Event Viewer
 - clearing filled logs, 806
 - described, 777–780, 805–806
 - S.M.A.R.T. messages from, 460
 - viewing audited information in, 1053–1054
 - Event Viewer snap-in in Computer Management, 775
 - exams
 - location of, 14–15
 - making arrangements to take, 14
 - required for the CompTIA A+ Certification, 10–13
 - testing process, 14–17
 - exceptions for a firewall, 1055
 - Exclusive Mode, 726–727
 - executable files, file extensions, 590
 - Executive Services, 517
 - Executive Software, defragmentation utility, 624
 - exhaust vents, 451
 - Expanded (EMS) memory, 723
 - expanded memory, 663, 667
 - Expanded Memory Manager (emm386.exe), 723
 - expand.exe utility, 784–785
 - expansion cards, 106, 663
 - expansion slots, 73–74, 96–97, 224. *See also* bus on the motherboard, 75
 - explorer.exe, 526, 993
 - Export Registry File dialog box, 770
 - extended attributes, 597–603
 - Extended Capabilities Port (ECP), 189, 246, 292
 - Extended Cylinder/Head/Sector (ECHS), 212–213
 - Extended Data Output (EDO) memory, 158, 166
 - Extended Industry Standard Architecture (EISA), 98–99
 - extended memory, 663
 - extended partition, 242–244, 760, 761
 - Extended (XMS) memory, 723
 - external cables, 306–308
 - external cache, 69
 - external drive, USB, 235
 - external knowledge bases, 26
 - external modem, 76, 323
 - external SCSI connectors, 228
 - external support professional, 27
 - external ventilation, 451
 - extract.exe, 834
-
- ## F
-
- F1 key, 184
 - F2 key, 183
 - F8 key
 - entering Safe Mode, 801
 - options during boot-up, 836–840
 - F10 keystroke, 183
 - fans
 - checking, 457
 - cleaning, 457
 - in a computer, 54
 - installing, 142–143
 - providing active exhaust, 452

- “Fast Ethernet” standards, 887
- Fast Page Mode (FPM), 157–158, 166
- Fast Pasting option, 727
- Fast ROM Emulation, 725
- Fast SCSI-2, 229
- Fast Wide SCSI-2, 229
- FAT (File Allocation Table), 250, 251, 261, 580
- FAT file system
 - cluster sizes, 237, 239–240
 - common to all Windows operating systems, 543
 - compared to NTFS and FAT32, 521
 - described, 237–238
 - NTFS compared to, 241
 - partitioning in, 243
- FAT12, 520
- FAT16, 520
- FAT32 file system
 - compared to FAT, 543
 - compared to NTFS, 241
 - compared to NTFS and FAT, 521
 - deciding to support, 760
 - described, 239–240
 - operating systems supporting, 240
- fault tolerance, 252–253
- `fc.exe` command, 622–623
- FDDI networking
 - architecture, 889
- `fdisk` command, 537
- `fdisk.exe` utility
 - described, 832, 835
 - preparing disks to hold data, 759–762
- feed rollers, 384
- feeder. *See* paper feeder mechanisms
- female connector, 298
- female port, 77
- fiber local connector (LC), 883
- fiber optic cabling, 297, 882–884
- File Allocation Table. *See* FAT
- File and Printer Sharing, 960–961
- file associations, 583–589
- file attributes, 594–607
- file compression, 591, 597–599
- file extensions, 586, 589, 589–594, 631
- file fragments, 765–766
- file sharing services, 939
- file systems
 - choosing, 520–522
 - optimization, 259
 - selecting in the Windows XP installation, 558
 - sharing resources, 959–968
 - specifying for a partition or volume, 251
 - supported by Linux, 506
 - supported by the Windows OS, 543
 - understanding, 237–241
- file tables, checking the structure of, 765
- File Transfer Protocol (FTP), 992–993
- file types, 584
- file virus, 1065–1066
- filenames, 581–583
- file-naming conventions, 579–594
- files
 - backing up, 23, 1046
 - breaking across a hard drive, 763
 - cleaning up temporary, 750
 - comparing the contents of two, 622–623
 - copying and moving, 620–622
 - deleting or removing, 621–622
 - encrypting, 599
 - moving, 621
 - not associated with an application, 584
 - not securable with FAT, 237
 - opening with `edit`, 628–629
 - renaming, 620, 622
 - restoring from a backup, 849–850
 - saving, 629
 - unintentionally deleting, 622
- Files and Settings Transfer (FAST) Wizard, 572–574
- filtering by MAC address, 915–916
- FilterKeys, 321
- Finder, 508
- `findramd.exe`, 835
- fingerprint reader, 1023
- finger-print scanners, 330
- Firefox, 994
- firewalls, 917, 1004–1005, 1020, 1054–1057
- FireWire, 51–52, 81–82, 293–294, 390–391
- FireWire 800, 82, 391
- FireWire connector, 305–306
- firmware, 54–55
 - defined, 152
 - for keyboards, 317
 - updates, 450, 918
 - upgrading on a printer, 391–392
- `fixboot`, 844
- fixed focus camera, 412

- fixmbr, 844
- flash cards, 236
- flash drives, 51, 235, 236
- Flash EPROM (Erasable Programmable ROM), 55
- flash RAM, 52
- flash ROM, 153, 181
- flashing, 181, 368
- flashlight, 472
- flat head screwdrivers, 36
- flat memory model, 647
- flatbed scanner, 408
- FlexATX motherboard, 95
- floppy disk
 - BIOS update program
 - from, 181
 - booting from, 191
 - described, 233–234
 - I/O address range of, 355
 - left in the floppy drive, 795, 796, 797
- floppy disk drive, 51
 - connector, 88, 276, 277–278
 - described, 85
 - enabling or disabling in CMOS, 187
 - maintenance of, 460
 - ribbon cable, 217
 - troubleshooting, 478
- flow control, 485
- Fn key on laptops, 491
- Folder Options dialog box, 586–587
- folders
 - connecting to shared, 966–968
 - encrypting, 599
 - sharing in Windows 2000, 961–964
 - sharing in Windows XP, 964–965
- Folders Explorer Bar, 525
- for loop in batch files, 636
- foreground, applications running in, 741

- form factors for PSUs, 279–280
- format command, 250, 252, 537
- format .com command, 626–627, 832
- formatting, 250–252
- formatting codes in HTML files, 991
- four-partition limit, 760
- FPM (Fast Page Mode), 157–158, 166
- FQDNs (fully qualified domain names), 950–954
- fragmented files, 624, 763
- frame type, 936–937
- free space
 - on drives, 748
 - required for defragmentation, 763
- friction feed printers, 380
- friendly names, resolving, 1000
- front side bus (FSB), 134
- FTP (File Transfer Protocol), 992–993
- F-Type connector, 879, 880
- full AT motherboard, 90–91
- full backup, 1030. *See also* normal backups
- Full Control permission, 963, 1049
- full redundancy, 253
- full-duplex communication, 898
- Fully Qualified Domain Names (FQDNs), 950–954
- function keys on a keyboard, 316
- fuser rollers, 381, 384, 401, 402
- fusing phase of laser printing, 384

G

- gateway, 895
- GDI (Graphic Device Interface), 711
- gdi32.dll, 661
- gdi.exe, 661
- gender changers, 308, 309
- General Protection Faults (GPFs), 514, 824
- Ghost, 539
 - .gif files, 785
- Gigabit Ethernet standards, 887
- gigahertz (GHz), 116
- glass cladding, 882
- Globally Unique Identifier (GUID), 194
- Gnome desktop, 504, 506
- Gopher protocol, 989
- goto statement, 634
- GPFs (General Protection Faults), 514, 824
- graphic files, 410, 593
- Graphic handling in Windows 9x, 661
- graphics accelerator cards, 337
- grounding of equipment, 456
- groups, 1040–1045
- guest account, 1041, 1042
- GUI (Graphical User Interface)
 - described, 321, 499–500
 - failure to start, 802
 - in Linux, 506–507
 - of the Macintosh OS, 508–509
 - of Windows, 502–503, 653
- GUID (Globally Unique Identifier), 194
- GUI-mode phase of Windows installation, 545, 559

H

- hackers, 1013–1014, 1015
- hal.dll file, missing, 798
- half-duplex communication, 898
- handheld devices, 438–439
- handheld scanner, 408
- handles, 738
- handshaking methods of WPA, 917
- handwriting recognition, 492, 729
- hard disks. *See* hard drives
- hard drive controller, 70, 213–214
- hard drives, 51. *See also* disk drives
 - access speeds of, 665
 - adding to laptop computers, 437–438
 - CMOS settings for, 185–187
 - configuring partitions on, 241
 - corruption of boot sectors on, 796
 - defragmenting, 460
 - in laptops, 423
 - limited to four partitions, 243
 - monitoring, 460
 - optimizing, 748
 - preventative maintenance for, 459–460
 - recognizing larger, 212
 - replacing on laptops, 437
 - sharing in Windows XP, 965
 - size of displayed in CMOS, 185, 186
 - terminology of, 207–213
 - troubleshooting, 478–480
- hardening systems, 1057
- hardware
 - defined, 54
 - detected by Windows 2000 setup, 550
- hardware address of a network card, 891
- hardware compatibility list (HCL), 541
- hardware devices, 683–688
- hardware level, reserving resources, 370
- hardware password. *See* power-on password
- hardware requirements for Windows, 540–541
- hardware resource settings, 685
- hardware tools, 35–38, 468–473
- hazardous materials, disposal of, 42
- HCL (hardware compatibility list), 541
- hcl.txt, 541
- HD DB-15. *See* DB-15 connector
- head alignment, 460, 478
- head positioning mechanism, 208, 209
- header of a network packet, 1054
- headphones, connecting, 343
- heat, 54
- heat sinks, 54, 142–143, 457
- Help and Support Center, 532
- help command, 628, 658
- help on the command line, 627–628
- hermaphrodite connector, 299, 300
- hertz, 116
- hex nuts, 471
- hexadecimal, 819
- hexadecimal addresses, 892
- Hidden attribute, 595, 596
- hidden devices, 800
- Hidden files, 605–606, 625
- hidden shares, 966
- hierarchical tree of Registry entries, 768
- high density disks, 478
- high memory, 664
- high memory manager, 657, 667
- High Performance File System (HPFS), 241
- high-voltage equipment, recognizing, 40
- himem.sys, 663, 664, 667, 723, 835
- hives, 772
- HKEY_CLASSES_ROOT, 767
- HKEY_CLASSES_USER, 767–768
- HKEY_CURRENT_CONFIG, 523, 768
- HKEY_CURRENT_USER, 769
- HKEY_LOCAL_MACHINE, 523, 647, 768, 769
- HKEY_USERS, 768
- HMA (High Memory Area), 664
- Hoaxbusters Web site, 1073
- home routers, 1004
- host adapter, SCSI, 223–224
- host ID portion of an IP address, 940
- HOSTS file, 952–953, 954
- hot fixes from Microsoft, 516
- hot-swappable technology
 - described, 435, 692
 - PCMCIA as, 102
 - SATA as, 232
 - USB as, 346
- HPFS (High Performance File System), 241
- HTML (HyperText Markup Language), 983, 990–992

HTT (hyperthreading technology), 121
HTTP (HyperText Transport Protocol), 989, 990, 1027
HTTPS, 992
hub, 81, 873, 896, 1059
humidifiers, 40
humidity, 428, 454, 455
hybrid topology, 874–875
Hypercard program, 989
hypertext, 989
hypotheses, 19, 20–21, 25

I

IAB (Internet Architecture Board), 981
IBM
 POST Beep Codes, 476
 Token Ring developed by, 888
 Type 1 connector, 299–300
IBM-type Data Connector.
 See IDC connector
IC (Integrated Circuit) chips, 54
icons
 representing file associations, 584
 specifying for MS-DOS applications, 719
 in this book, 4–5
IDC connector, 888–889
IDE (Integrated Device Electronics)
 40-pin connector, 87
 cabling, 217–219
 chain, 219
 compared to EIDE, 214
 connections on motherboards, 85–86, 298
 controllers. *see* IDE connectors
 described, 51
 devices, 213–222
 drives, 85, 217, 218
 ribbon cable, 85–86, 217, 218
 standard, 215
IDE/ATA standard, 215
Idle Sensitivity setting, 727
IEEE 1284 standard, 292, 293, 306, 346, 388
IEEE 1394. *See* FireWire
IEEE 1394b standard, 82
IEEE standards, 889–891
if statement, 633–634
illegal operation, 824
images, scanning, 409–411
imaging software, 539
IMAP (Internet Message Access Protocol), 988–989
impact printers, 386
incompatibilities, 715–716, 727–728
incremental backups, 596, 625, 848, 1030
Indexing Service, 602–603
Industry Standard Architecture. *See* ISA
 .inf files, 682
 .info file, 777
infrared devices, 390, 907–908
infrared ports, 360, 390, 900
infrared printer connections, 390
infrared wireless communication, 432–433
infrastructure mode for wireless, 909
inheritance option, turning off for permissions, 1050
 .ini files, 522–523, 654, 770, 771
 structure of, 671
initial cache, 670, 671
initialization vector (IV), 916
inkjet printers, 385–386, 401, 462–463
inline UPS, 283
input devices, 56–59, 315–330, 429–431
input voltage range, 273
input/output (I/O) ports, 74, 75
installation
 preparing for, 540–544
 verifying and troubleshooting
 Windows, 564–565
installation files, updating, 571–572
instances in System Monitor, 744
instruction pipeline, 129
INT13 extensions, 212
integrated cache, 69
integrated sound cards, 344
integrity of data stored in memory, 164
Intel processors, 127–136
interference for wireless networks, 920
internal cables, 298
“internal cache”, 165
internal components, placement of, 454
internal modems, 323, 324–325
internal ribbon cables, 298
internal SCSI device, 226
internal support professional, 27
internal ventilation, 452–454
Internet
 accessing, 1002–1004
 history of, 981–983
 as a large hostile remote network, 917

protocols, 985–994
 resources, 26
 using, 998–1001
 Internet Architecture Board (IAB), 981
 Internet Explorer, 994, 995–998
 Internet Message Access Protocol (IMAP), 988–989
 Internet Options dialog box, 995–998
 Internet Protocol Security (IPSec), 918, 1028
 Internet service providers (ISPs), 983–985
 Internet Services section in Network Diagnostics, 780
 Internet Zone sites, 995, 996
 Interrupt Requests. *See* IRQs
 intrusion detection option in CMOS, 1040
 invalid working directory, 824–825
 inventory list, CMOS as, 179–180
 invitation file from Remote Assistance, 532
 I/O addresses
 configuring, 189, 365
 described, 352–355
 dynamically assigned, 101
 not required for SCSI devices, 224
 for serial ports, 190
 of standard ports, 354
 I/O devices, 346
 io.sys, 646, 653–654, 833
 IP, 986
 IP addresses
 comparing with subnet mask, 943
 configuring, 940–942, 944
 converting to MAC, 953
 pinging, 956
 resolving computer names to, 949
 unique for each computer, 1000
 used by WINS server, 949
 viewing, 931
 ipconfig command, 946
 ipconfig utility, 955, 956
 IPsec (Internet Protocol Security), 918, 1028
 IPX/SPX protocol, 936–937, 986
 IrDA (Infrared Data Association), 907
 IRQ 5, 344
 IRQs (Interrupt Requests)
 assigning with jumpers, 89
 cascading, 357–360
 configuring, 189, 365
 dynamically assigned, 101
 for each serial port, 190
 generally available, 360
 line, 355–360
 not required for SCSI devices, 224
 standard, 358, 359
 troubleshooting, 366
 ISA (Industry Standard Architecture)
 bus, 56, 97
 devices, 367
 network card, 97, 98
 non-Plug and Play devices, 198
 slots on the motherboard, 74, 75
 sound cards, 344
 ISDN (Integrated Services Digital Network), 58, 1003
 Isopoint pointing device, 491
 ISPs (Internet service providers), 983–985
 Itanium processor, 134–135

J

JavaScript, 783–784
 journal-based file systems, 506
 journaling file system, NTFS as, 521
 JPEG file, 410
 jumpers
 arranging the SCSI ID, 226
 configuring resources, 367–368
 described, 88–89
 removing CMOS passwords, 192, 193
 removing to write to the BIOS chip, 183
 setting for hard drives, 478, 479
 settings on IDE hard drives, 220–221

K

K6 processors, 136
 KDE desktop, 504, 506
 kernel memory, 738, 818
 kernel mode, 517, 518
 kernel1132.dll, 661
 kernel-mode processes, 517
 key pairs in PKI, 599
 key words in exam questions, 15
 keyboard(s)
 configuring, 318–321
 disassembling, 481, 482
 groupings on, 315–316
 installing, 318
 laptop, 423, 429–430
 preventative maintenance, 459
 troubleshooting, 481
 types of, 317–318
 using comfortable, 45

keyboard connectors, 294
keyboard controller, 70
keyboard device driver, 317
keyboard port, 59
keyboard/mouse connectors, 78–79
keycaps, removing, 481
keycode, 316
keyed cables, 298
keyed connectors, 295, 298
keyed DIMMs, 168
keyed power connectors, 83
keys
 per-session in WPA, 916
 in the Registry tree, 768
keyswitch on each keyboard
 key, 316
knife, 473
Knoppix, 506, 856–859
Knowledge Base, links to,
 779
knowledge bases, 25, 26
krn1386.exe, 661

L

L1 (level-1) cache, 69, 119,
 128, 131, 165
L2 (level-2) cache
 described, 69, 70, 119, 165
 in a dual core processor,
 121
 increasing performance,
 106
 on the Pentium Pro, 130
L3 cache, 135, 166
label command, 595
:label statement, 633–634
LAN (Local Area Network),
 871, 889–891, 931, 932
LANguard Network Security
 Scanner, 1077
laplink cable, 899

laptop computers
 AC adapters for, 280–281
 benefits of, 421
 changing batteries on,
 425–426
 communication
 components of,
 432–433
 described, 421–423
 expanding, 434–436
 including wireless
 network cards, 909
 input devices, 429–431
 installing memory on,
 169–171
 maintenance of, 461–462
 power management
 functions of, 426
 projecting images from,
 341
 troubleshooting, 489–492
 upgrading, 436–438
large disk support, 760
laser printers, 380–384, 401,
 402–403
laser printing process,
 383–384
laser, write operation with,
 235
Last Known Good Configu-
 ration, 695, 817, 839,
 850
latency, 211
layered approach to
 troubleshooting, 22–23
layered network model, 23
LBA (Large Block Access)
 mode, 479
LBA (Logical Block
 Addressing), 212–213
LBAenabled BIOS, 212
LC (fiber local connector),
 883

LC4 password-cracking
 tool, 1016
LCD (Liquid Crystal
 Display)
 cutoff switch, 491
 described, 340
 monitors, 294
 panels, 423, 427–429
 screens, 428, 429, 461
Least Recently Used (LRU)
 algorithm, 666
LEDs on a keyboard, 316
legacy devices, 198,
 692–693, 826
lfnbk.exe, 583
LFX v1.1 (Lower Profile
 Form Factor), 279, 280
lifespan of equipment, 457
Li-Ion batteries, 423, 424,
 462
LIM EMS (Lotus Intel
 Microsoft Expanded
 Memory Specification),
 663
limits in computers, 580
Lindner, Paul, 989
line conditioners, 282
linear format exams, 10
line-in port, 79, 80, 343
line-of-sight infrared
 devices, 390, 433,
 907–908
line-out port, 343
lines on the screen for MS-
 DOS applications, 724
link speed of a network
 card, 328
Linux, 503–507, 1005
Li-Poly batteries, 461, 462
liquid cooling system for a
 processor, 142
Liquid Crystal Display. *See*
 LCD

liquids as problems, 454
 listening ports, 958
 listsvc command, 844
 lithium-based batteries, 462
 lithium-ion polymer (Li-Poly) batteries, 461, 462
 Live CDs, 38, 505, 856
 LMHOSTS file, 949–950, 954
 Load and Unload Device Drivers Properties dialog box, 693
 loadstate, 574
 local area connection icons, 930, 931
 Local Area Network. *See* LAN
 local Intranet sites, 995, 996
 local logon, 1046
 local printer in Windows, 392
 local security policies, 1047–1048
 Local Security Policy application, 693
 Local Security Policy console, 1052
 Local Users and Group snap-in, 775
 lock icon in the Web browser, 1027
 lockdown cable, 462, 1022
 lockups, 819–820
 Log On tab, 753
 logging and reporting tool, 777–780
 logic bombs, 1020
 Logical Block Addressing (LBA), 212–213
 logical drives, 242, 243, 760, 761
 logon command, 844
 logon type, 556
 log-on-locally right, 1046

logs, 564, 744–745, 849
 long filenames, 582, 583
 loopback address, 941–942, 956
 loopback plugs, 473
 looping in batch files, 636–637
 lost chains, 765–766
 Low Voltage Differential. *See* LVD
 LPT1, 76, 355, 360
 LPT2, 355, 360
 LPX (low profile extended) motherboard, 93
 LVD (Ultra2 SCSI), 230

M

MAC addresses, 891–892, 915–916, 931, 953
 Macintosh OS, 502, 507–509
 Macintosh OS X, 505, 507, 584
 Macintosh proprietary hardware, 508
 macro security settings, 1066–1067
 macro virus, 1066
 macros, malicious, 1066
 magnetic pickup tools, 470
 magnetic screwdrivers, 36
 magnetic tape drives, 51
 main power connector, 277
 maintenance kits for laser printers, 463
 maintenance tasks, 457–463
 makeboot.bat, 538
 male Centronics connector, 306, 387
 male connector, 298
 male DB-25 connector, 306, 387
 male port, 77

malicious software, 1065–1073
 malware, 826
 management tools, 259–263
 man-in-the-middle attack, 1018
 manufacturers, hardware devices, 683, 684
 mapping network drives, 967–968
 Master Boot Record (MBR), 196, 212, 762
 Master File Table (MFT), 580, 598
 master/slave chain, 215, 216
 master/slave configuration for IDE hard drives, 219–222
 Material Safety Data Sheet (MSDS), 43
 math coprocessor, 119–120, 355
 MAU (Multistation Access Unit), 888
 /mbr switch of fdisk.exe, 762
 MBSA (Microsoft Baseline Security Analyzer), 1076–1077
 MCA architecture, 98, 99
 MDA standard, 336
 MDRAM (Multibank DRAM), 338
 Mean Time Before Failure (MTBF), 457
 mechanical mouse, 321
 Media Access Control address. *See* MAC addresses
 media problems, 481
 megahertz (MHz), 50, 116, 154
 megapixels, 412, 417

- mem command, 624, 668
- membrane keyboard, 316
- memory
 - accessing data in, 157
 - addresses, 363
 - amount indicated in
 - CMOS, 187–188
 - causing auto restarts, 816
 - conventional, 662–663
 - defined, 151
 - description of, 52–53
 - determining the type of, 166
 - expanded, 663
 - extended, 663
 - installing on desktop PCs, 168–169
 - installing on laptop systems, 169–171
 - installing or upgrading, 166–171
 - in laptops, 423
 - for MS-DOS applications, 721
 - organization of, 156–157
 - paging, 518–520
 - PDAs storing information in, 438
 - processor accessing, 116
 - purpose of, 151–152
 - troubleshooting, 478
 - types of, 151–156
 - upgrading in laptop computers, 436
 - upgrading on printers, 391
 - in Windows 9x, 662–669
- memory bank, 160
- memory controller, 70, 118, 158
- memory leak, driver with, 738
- memory management, 514
- memory manager, 663
- memory modules, 69, 70, 167, 168
- memory packages, 160–163
- memory slots, 68, 160
- memory space, 707, 711, 822–823
- memory sticks. *See* flash drives
- memory usage in Task Manager, 736, 737
- memory.dmp file, 520
- menu bar in the Mac OS, 508, 509
- mercury batteries, 43
- message queue, 709, 711
- Messenger service, 971–972
- Metropolitan Vacuum Cleaner, 449
- MFT (Master File Table), 580, 598
- mic port, 416
- mice. *See also* mouse
 - installing, 322–323
 - preventative maintenance, 459
 - troubleshooting, 481
 - types of, 321–322
- MIC-in port, 79, 80
- Micro ATX power supplies, 274
- Micro Channel Architecture (MCA), 98, 99
- MicroATX motherboard, 95
- MicroDIMM, 162
- Micro-PGA packaging, 138
- microphone(s), 343, 416
- Microsoft
 - malicious code removal tool, 1072–1073
 - testing device drivers, 689
- Microsoft Baseline Security Analyzer (MBSA), 1076–1077
- Microsoft Management Console (MMC), 529–530, 775. *See also* Services MMC snap-in
- Microsoft Network Monitor, 1017–1018
- Microsoft Office, 590–591, 1066–1067
- Microsoft Update, 817, 1057, 1074
- microwaves, affecting wireless networks, 921
- MIDI (Musical Instrument Digital Interface), 310, 416–417
- Millions of Instructions Per Second (MIPS), 50
- MIME (Multipurpose Internet Mail Extensions), 987
- mini dump files, 781
- Mini-DIN 6 connector, 294, 318, 319. *See also* PS/2
- mirror, small, 471
- mirrored volumes, 247, 1031
- mirroring. *See* disk mirroring
- “Missing Operating System” message, 24
- MJPEG compression, 415
- mkdir (md) command, 619
- MMC. *See* Microsoft Management Console; multimedia cards
- mmc.exe, 530
- MMF (multimode fiber), 883
- MIMO (Multiple-Input Multiple-Output) technology, 913
- MMX (MultiMedia eXtensions), 120
- Mobile Duron, 136
- mobile version of processors, 135–136
- modem commands. *See* AT commands

- modems
 - characteristics of, 323
 - configuring, 325–326
 - described, 58
 - installing, 324–325
 - troubleshooting, 484
 - verifying connection speed and data characteristics, 485
- Modems and Network Adapters section, 780
- Modify permission, 1049
- modulation, 58, 323
- molex connectors, 276, 278, 279, 298
- monitor adjustment tools, 472
- monitor connector, 294
- monitoring tools, 736–746
- monitors. *See also* CRT monitors; displays, types of; LCD (Liquid Crystal Display)
 - cleaning, 458
 - common symptoms, 483
 - external causing issues with laptops, 491
 - installing an additional, 342–343
 - keeping at eye level, 45
 - maintenance of, 41–42
 - maintenance tasks, 458
 - overview of, 56–57
 - troubleshooting, 482–484
 - using, 339–343
- Monochrome Display Adapter (MDA), 336
- monochrome LCDs, 428
- Moore's Law, 910
- Morris Worm, 982
- motherboard(s)
 - cache residing on, 166
 - chipset, 70–71
 - controlling the power supply, 273
 - described, 53, 65
 - fitting into cases, 96
 - form factors, 89–90
 - identifying the types of, 89–96
 - performance considerations, 106
 - speed of, 128, 159
 - troubleshooting, 477
- motion sensor for a laptop, 462
- mounting kits for 3.5-inch drives, 473
- mouse. *See also* mice
 - using comfortable, 45
- mouse and keyboard connectors, 78–79
- mouse ball, 430
- mouse port, 59
- move command, 621
- Mozilla, 994
- MPEG (Moving Picture Experts Group), 415
- MRTG tool, 460
- msconfig.exe, 766
- MS-DOS
 - applications, 518, 726
 - boot disk, 793, 794
 - boot process for, 653–662
 - described, 500
 - disk caching utility, 669
 - file system created by, 580
 - files from Windows 98, 653
- MS-DOS Mode, 840
- MS-DOS-based applications, 667, 716–728
- MS-DOS-based disk utilities, 582
- msdos.sys, 654, 654–656, 833
- MS-DOS/Windows 9x, boot loader, 653
- MSDS (Material Safety Data Sheet), 43
- .msi files, 700
- mstsc.exe, 531
- MTBF (Mean Time Before Failure), 457
- Multibank DRAM (MDRAM), 338
- multi-display support, 342–343
- Multifunction controller, 649
- multihead screwdriver, 36
- multimedia, 310
- multimedia cards (MMC), 346
- multimedia devices, 407–417
- MultiMedia eXtensions (MMX), 120
- multimeter, 37, 468–469
- multimode fiber (MMF), 883
- multiple access in CSMA/CD, 884
- multiple branch prediction, 130
- multiple shares, 966
- Multiple-Input Multiple-Output (MIMO) technology, 913
- multiprocessing system, 133
- multiprocessor computer, 706
- Multipurpose Internet Mail Extensions (MIME), 987
- Multistation Access Unit (MAU), 888
- multitasking, 120, 704–706
- multithreaded applications, 121, 706, 738
- multithreaded
 - environment, 705
- Musical Instrument Digital Interface (MIDI), 310, 416–417
- My Computer, 524–525, 526
- My Network Places, 526

N

- name resolution, 947–954
- nanoseconds (ns), 52, 154
- NAT (Network Address Translation) gateways, 1004
- National Center for Supercomputing Applications (NCSA), 990
- natural keyboard, 317–318
- navigating computers, 523–527
- nbtstat utility, 957
- NCSA Mosaic, 990
- needle-nosed pliers, 470
- net command, 812, 813
- NetBEUI protocol, 936
- NetBIOS name cache, 957
- NetBIOS name resolution, 948, 957
- NetBIOS names, 947–948.
See also computer names
- NetBIOS over TCP/IP (NetBT), 987
- NetBIOS over TCP/IP Statistics, 957
- NetBIOS session interface, 987
- NetBus, 1019, 1071
- Netscape Navigator, 994
- netstat command line utility, 958
- network(s)
 - access methods, 884–886
 - adding programs from, 703
 - administrative credentials for, 29
 - architecture, 886–889
 - attacks causing auto restarts, 816–817
 - booting off, 196
 - browsing resources, 966–967
 - cabling, 876–884
 - client , 932–935
 - components, 939
 - defined, 867
 - devices, 891–898, 956
 - disabling temporarily, 931
 - drives, mapping, 967–968
 - filtering data, 896
 - hub, 896
 - protocols, 889–891
 - switches, 896–897
 - topologies, 871–876
 - traffic, 894, 1017–1018
 - types of, 867–871
- network adapter driver, 928–932
- network adapters (cards), 79, 80, 891–892
 - built-in, 196
 - configuring, 327–328
 - installing, 326–327, 929
 - types of, 928
 - using to network two computers, 899
 - working with, 326–328
- Network connectivity, 927
- Network Diagnostics tool, 780
- network ID in an IP address, 940
- Network Identification Wizard, 555–556
- network installation of Windows, 539
- network interface cards (NICs). *See* network adapters (cards)
- Network Neighborhood. *See* My Network Places
- Network Operating System (NOS), 869, 871
- network packets, parts of, 1054
- network ports, disabling, 1021
- network printer, 392, 396–397, 968–970
- network service, starting up, 652
- network sniffer on wireless networks, 916
- Network Stumbler, 921, 922
- Network+ certification, 5
- network-based attacks, 1015–1018
- network-based printers, 388–389
- networking
 - cables, 884
 - computers, 898–900
 - configuring for Windows 2000, 553, 554
 - software components of, 927–939
 - tasks for troubleshooting, 931–932
- Networking tab of Task Manager, 714, 738–739
- New Partition Wizard, 245–246
- New Technology File System. *See* NTFS
- New User dialog box, 1042–1043
- New Volume Wizard, 249, 254–256, 258–259
- NIC (network interface card), 79, 80
- NiCad battery, 423
- Nickel-Metal Hydride batteries, 423–424
- Nimda virus, 1072
- NLX (New-Low-profile extended) motherboard, 93
- “No operating system found” error, 795, 796

- noise words in the indexing service, 602
- non-cooperative programs, 705
- non-paged memory, 738
- non-parity memory, 163–164, 167
- non-Plug and Play devices, 692–693
- resources, 367–368
- non-reentrant programming code, 710
- non-routable protocol, 938
- non-terminated end in a bus topology, 872
- nonunicast traffic, 739
- normal backups, 848
- normal.dot, 590
- North Bridge chip, 70, 71
- Norton AntiVirus, 1068–1070
- NOS (Network Operating System), 869, 871
- Notepad files, 588
- notepad.exe, changing config.sys, 794
- Notification Area, 530
- Novell networks, connecting to, 895
- Novell server, matching the frame type of, 937
- nslookup TCP/IP utility, 958
- ntbackup.exe, 450, 845
- ntbootdd.sys file, 647, 648, 797, 799
- ntbtlog.txt log file, 839
- ntdetect.com, 647, 650, 798
- NTFS (New Technology File System) benefits of using, 521–522 compared to FAT or FAT32, 543 described, 240–241, 500 partition formatted for, 1049 partitions, 607 permissions, 1048
- ntldr, 544, 646, 647, 797 “NTLDR not found” message, 24
- ntoskrnl.exe, 647, 650–651, 796, 798–799
- ntuser.dat, 773, 801
- NTVDM (NT Virtual DOS Machine), 518, 712, 820
- null modem adapter, 309–310
- null modem cable, 307–308, 899, 900
- number lock on a laptop, 491
- number pad functions on laptops, 491, 492
- numeric keypad on a keyboard, 316 securing the entrance to a facility, 1020, 1021
- Numeric Processing Unit (NPU). *See* math coprocessor
- nut drivers, 471
- NWLink (IPX/SPX), 936, 937
-
- ## O
-
- oakcdrom.sys, 834
- object access, auditing for, 1052
- obvious, not discounting the, 24
- octets in an IP address, 940
- odd parity, 163, 164, 485
- offline UPS, 283
- ohms, 272
- OLGA (Organic Land Grid Array), 134
- on-board cache, viewing, 197–198
- online UPS, 283
- Open dialog window, 629
- Open With dialog box, 585
- operating system(s). *See also* OS kernel boot code on the MBR, 212 checking the version of, 514–517 comparison of, 499–509 components of, 517–523 functions of, 513–532 identifying major functions, 513–517 rights, 848 supporting Plug and Play, 690 valid licenses for, 540 operating systems section in boot.ini, 648
- “opposites attract” principle, 384
- Opteron, 137–138
- optical drives, 51, 480
- optical fiber, 882
- optical media, repairing or polishing, 461
- optical mouse, 322
- optical zoom, 412
- optimization, best practices for, 746–755
- optional devices, not functioning, 825–826
- opto-mechanical mouse, 321–322
- Organic Land Grid Array (OLGA), 134
- organizational units in Netware, 934
- original location, restoring to, 849
- OS. *See* operating system(s)

OS kernel
of Linux, 503
loading in Windows 9x, 660
for Windows XP, 650
OS X. *See* Macintosh OS X
OS/2, 241, 657
OS-level
software modems, 484
tools for maintenance, 450
outcomes, documenting, 25
output devices, 335–346
overclocking, 122
ownership of files and other objects, 1046

p

P4 connector, 84, 278
PAC (Pin Array Cartridge)
packaging, 134
packaging of Linux distributions, 504
Packet Internet Groper. *See* ping.exe utility
packet sniffers, 1017
packets, 1017
page faults, 740
Page file usage, 737–738
page pool, 738
page printers. *See* laser printers
paged memory, 738
pagefile.sys, 746
Pages/sec counter, 742
paging file
configuring, 519–520
described, 723
errors, 801–802
using, 665
in Windows XP, 746–748
paging memory, 518–520
pairing Bluetooth devices, 919–920

PAN (Personal Area Network), 919
paper feeder mechanisms, 379–380, 381
paper jams, 398, 400, 401
parallel cable, 306, 387, 388
parallel devices, 346, 484
parallel port mode, 189, 409
parallel ports
configuring in CMOS, 188–189
connecting scanners, 409
described, 59, 292
disabled in the Device Manager, 485
on the motherboard, 76–77
networking two computers, 899–900
in a printer, 387
troubleshooting, 484–485
parallel processing, 129
parameters
in batch files, 632–633
incorrect for printers, 814
parity, 485
parity bit, 164
parity data, 247, 1032
parity information, 257
parity memory versus non-parity, 163–164, 167
partition(s)
configuring on a disk, 241
creating on a disk, 760
defined, 210, 242
deleting existing, 542
formatting, 250–252
managing, 241–252
obtaining an NTFS, 1049
specifying the format of, 627
during Windows 2000 installation, 549
in Windows 2000/XP/2003, 244–249

in the Windows OS, 542–543
in the Windows XP installation, 558
partition number in an ARC path, 649
partition table, 212, 761
passive heat sink, 142
passive matrix display, 428
password attacks, 1015–1016
passwords
changing for wireless networks, 915
checking weak, 1077
policies for, 1058
setting in CMOS, 192
strong, 1016, 1025, 1026
for user accounts, 1042
patches, 1074, 1075
patch.exe file, 1019, 1071
patching systems, 1057
path environment variable, 637
pathping command line utility, 958
PC cards. *See* PCMCIA cards
PC repair toolkit, 35–36
PCI (Peripheral Component Interconnect)
bus, 51, 56, 101
card with SATA connectors, 232
described, 104–105
modems shipping as cards, 324
network adapter, 928
network card, 326, 327
slots compared to ISA and EISA, 224
slots on the motherboard, 74, 75
sound cards, 344

- PCIe (Peripheral Component Interconnect Express) bus, 57
- PCI-X (Peripheral Component Interconnect Extended) bus, 57, 103
- PCMCIA, 102–103
- PCMCIA cards, 102, 422, 434–436
- PCMCIA (PC Card) buses, 52
- PCMCIA slot, 391
- PDAs (Personal Digital Assistants), 438–439
- peer servers, 869
- peer-to-peer (P2P) networks, 867–869
- Pentium 4, 67, 127, 134, 135
- Pentium chips, 67, 119
- Pentium II, 66, 67, 127, 131, 133
- Pentium III, 67, 127, 132–133
- Pentium III M, 136
- Pentium III Xeon chip, 133
- Pentium “M”, 135–136
- Pentium Pro processor, 130
- Pentium processor, 127–129
- perfmom.exe, 741
- performance
 - of a hard drive, 217
 - increasing for processors, 143–144
 - of motherboards, 106
 - of striped volumes, 247
 - tuning at the command prompt, 623–624
- Performance Logs and Alerts snap-in, 775
- Performance Monitor, 480, 741–746
- Performance objects for critical system resources, 743–744
- Performance tab of Task Manager, 714, 736–738
- Peripheral Component Interconnect. *See* PCI
- peripheral connector, 276, 278
- permissions
 - assigning to particular users, 960
 - available, 1049
 - configuring in NTFS, 240
 - described, 1048–1051
 - implementing, 1045–1046, 1048–1051
 - inheritance of, 1049, 1050
 - to install device drivers, 693–694
 - on Registry keys, 773
 - setting on a resource, 1026
 - setting on a share, 962
 - Windows XP errors
 - related to, 729–730
- Personal Area Network (PAN), 919
- Personal Computer Memory Card Industry Association. *See* PCMCIA
- personal computers (PCs)
 - disposing of components, 42–43
 - functions of, 50
 - installing memory on desktop, 168–169
- Personal Digital Assistants (PDAs), 438–439
- PF (Page File) in Task Manager, 736, 737
- PF (Power Factor), 283
- PFC (Power Factor Converted) power supply, 284
- .PFX file, 601
- PGA (Pin Grid Array) chip, 123, 124, 125
- Phillips screwdrivers, 36
- phone line issues, 484
- Photomultiplier Tube (PMT), 408
- photos, scanning, 410
- photosensitive drum, 383
- physical access, 1039
- physical contact, with media surface, 461
- physical environment for troubleshooting, 475
- physical memory, 514, 710, 738
- physical RAM, 723
- physical security, 1020–1023
- pickup rollers, 380, 462
- pickup tools, 470
- pictures, uploading, 413
- .pif files, 717–727
- .pif Properties window, 718–727
- PIFs (Program Information Files), 822
- Pin 1
 - on a DIMM, 168
 - on a DIP chip, 123
 - for an IDE hard drive, 219
 - on a PGA, 123
 - on a processor chip, 140
 - on a SIMM, 168, 169
- Pin 1 rule, 86, 218
- Pin 40 on an IDE hard drive, 219
- Pin Array Cartridge (PAC) packaging, 134
- Pin Grid Array (PGA) chip, 123, 124, 125
- PIN (Personal Identification Number), 1025
- ping of death, 1017
- ping utility, 802
- ping.exe utility, 955–956

- PIO (Programmed I/O) mode, 217
- pipe-more command, 618
- PIV (Pentium 4), 133
- pixels, 337, 411, 483
- PKI (Public Key Infrastructure), 599
- planning for installation, 540–544
- platen, 408, 463
- Platform Support Reporting Utility, 635
- platters, 208–209, 210
- Plug and Play
 - architecture, 101
 - BIOS (PnP BIOS), 193
 - described, 193, 690–693
 - process, 690–691
 - requirements for a full environment, 369
 - resources, 368–369
 - technology, 102, 389
- PMT (Photomultiplier Tube), 408
- PnP BIOS (Plug and Play BIOS), 646
- Point and Print, 969–970
- polarity base plugs, 456
- POP3 (Post Office Protocol version 3), 988
- pop-ups, blocking, 997
- port numbers, 1055
- port replicators, 433
- portable computers. *See* laptop computers
- ports
 - with associated addresses, 360
 - choosing for printers, 814
 - connecting a scanner, 409
 - defined, 291
 - disabling unused on a switch, 1059
 - identifying common, 291–294
 - on the motherboard, 74–82
 - overview of, 58–59
- Post Office Protocol version 3, 988
- POST (Power-On Self Test)
 - contained in BIOS, 180
 - described, 72
 - errors, 475–476
 - process, 646
 - stored in ROM, 153
- power
 - adapters, 281, 490
 - basics of, 271–272
 - conditioning, 456–457
 - connectors, 82–84, 276–279
 - consumption, 456
 - converters, 271
 - cords, 488
 - interruptions, 456
 - issues with, 282–284
 - management of laptop computers, 426
 - troubleshooting problems, 37
- Power Factor Corrected (PFC) power supply, 284
- Power Factor (PF), 283
- power supplies
 - causing auto restarts, 816
 - connecting to an IDE hard drive, 219
 - danger in servicing, 41
 - described, 53, 271
 - exceeding the total output of, 488
 - form factors, 279–280
 - identifying, 274–276
 - in laser printers, 382
 - maintenance tasks
 - associated with, 458
 - motherboard controlling, 273
 - result of defective, 489
 - tasks related to, 272–273
 - troubleshooting, 487–489
- power supply tester, 488
- power supply units (PSU), 41, 272–273, 279–280
- power surges, 282
- Power Users group, 1044
- power-on password, 192, 1040
- Power-On Self-Test. *See* POST
- preemptive multitasking, 705–706
- pre-logon default user, 773
- pre-shared key (PSK), 917
- preventative maintenance, 402–403, 447–463
- pre-Windows 2000 applications, 701–702
- primary controller for EIDE, 215
- primary corona wire, 382
- primary hard disk controller, 355
- primary partition, 242, 761
- primary ring in FDDI, 889
- primary-rate ISDN, 1003
- print device, 392
- print drum, 382, 383
- print head, 386–387
- Print Operators group, 1044
- print queue, 812
- print servers, 389
- print spooler, 812
- print spooler folder, 399–400
- print spooler service, 399, 971
- printer(s)
 - configuring in Windows, 394–396
 - correcting slow, 398
 - defined in Windows, 392

- incorrect parameters for, 814
 - incorrect/incompatible driver for, 813–814
 - installing, 392–394
 - maintenance of, 462–463
 - managing, 379–403
 - sharing, 395–396, 961
 - troubleshooting, 397–401
 - types of, 380–387
 - upgrading, 391–392
 - printer cable, 77
 - printer drivers, 392, 813–814, 969
 - printer port, 76
 - printer resources, sharing, 939, 968–970
 - printing problems, Windows-specific, 811–814
 - printing process, optimizing, 748–749
 - priorities, 394, 395, 740
 - privacy, 29, 996–997
 - private keys, 599
 - privilege use, auditing, 1052
 - probe leads in a multimeter, 469
 - problems
 - managing common, 811–826
 - not belittling or minimizing, 30
 - preventing, 1074–1078
 - prioritizing, 21–22
 - process tracking, auditing, 1052
 - process tree in Task Manager, 740
 - processes in Task Manager, 736, 737, 738
 - Processes tab of Task Manager, 714, 739–740, 820
 - processors
 - AMD, 136–138
 - characteristics of, 115–125
 - CPU, 50–51
 - identifying on the motherboard, 66–67
 - increasing performance, 143–144
 - installing, 138–141
 - Intel, 127–136
 - keeping cool, 141–143
 - overclocking, 122
 - retrieving information from memory, 165
 - sockets, 125, 126, 138
 - speed of, 50, 116
 - temperature checking, 142, 457
 - troubleshooting, 477
 - product key, 552
 - professional behavior, 28–31
 - program information files. *See* .pif files
 - Programmed I/O (PIO) mode, 217
 - programs, 700
 - programs settings in Internet Explorer, 997–998
 - projectors, 340–341
 - Protected Mode memory, 723
 - protected-mode processors, 120
 - protecting yourself and your computer, 35–45
 - protocol suites, 938, 986
 - protocols, 935–938, 985
 - proxy servers, 997, 998
 - PS/2
 - connection, 317
 - connector, 78, 79, 304, 305, 318, 319, 323
 - described, 294
 - mouse, 323
 - PSK (pre-shared key), 917
 - PSU (power supply units), 41, 272–273
 - public access terminals, 983
 - Public Key Infrastructure (PKI), 599
 - public keys, 599
-
- ## Q
-
- QIC (Quarter Inch Cartridge) standard, 237
 - QTParted, 858, 859
 - Query Modem button, 325
 - Quick Edit for MS-DOS applications, 726
 - quick format, 252, 626
-
- ## R
-
- Radio Frequency Interference (RFI). *See* EMI
 - Radio Grade (RG), 879
 - RAID, 252–259, 1031–1032
 - RAID 5 volume, 247, 257, 1032. *See also* RAID level 5
 - RAID level 0, 253–256, 1031
 - RAID level 1, 253–256, 1031
 - RAID level 5, 253, 257–259, 1032
 - RAM (Random Access Memory)
 - chips relying on CMOS, 55
 - clean memory location in, 666
 - described, 52, 153–156
 - drive, 834
 - PDA's storing information in, 438
 - requiring power to retain information, 52
 - types of, 154–156

- Rambus DRAM (RDRAM), 159
- Rambus Inline Memory Module (RIMM), 159
- ramdrive.sys, 835
- RCA cables, 310
- RD (receive data) wire, 900
- RD (Received Data), 307
- rdisk.exe utility, 841
- RDP (Remote Desktop Protocol), 530, 531
- RDRAM (Rambus DRAM), 159
- read caching, 669
- Read permissions, 964, 1049
- Readiness Analyzer, 566
- Read-only attribute, 595, 596
- read-only files, 625
- Read-Only Memory. *See* ROM
- read/write head, 208, 209
- read/write process, 210–211
- Ready To Send (RTS), 307
- real-mode, 653, 654
- real-mode processor, 120
- real-time protection, 1068, 1069–1070
- rebooting the computer, 711
- Received Data (RD), 307
- receptacle tester, 487
- rechargeable batteries, disposal of, 43
- recovery agent in EFS, 240–241, 600, 601
- recovery CDs, 856
- Recovery Console, 38, 546, 843–844
- recovery partitions, 856
- Recovery tab, 753–754
- Red Hat Package Manager (RPM), 504–505
- redirection symbols, 632
- redundancy, 252, 253
- Redundant Array of Inexpensive Disks. *See* RAID
- reflective infrared devices, 390
- refresh rate, 340
- .reg files, 771, 772
- regedit.exe, 833
- regedt32.exe, 770, 772–774
- reg.exe, 773
- registered files types, 589
- registers, 118
- registration rollers, 381. *See also* feed rollers
- Registries, connecting to remote, 772
- Registry
 - backing up, 846
 - categories, 523
 - corruption of, 800–801
 - device and service settings in, 651–652
 - editing, 585
 - in the Emergency Repair Process, 842–843
 - export files, 770
 - exporting sections of, 770
 - hives, 772
 - making a clean backup of, 769
 - referencing devices, 800
 - role in Plug and Play, 691, 692
 - settings, 973
 - sub-trees of, 767–768
 - UnInstall key, 703
 - user portion of, 801
 - using care when working in, 704
 - viewing file associations, 586
 - of Windows, 522–523
 - in Windows 9x, 661
- Registry Editor, 768–770, 772
- Registry keys, 772, 773
- re-installation, 24
- relative path, 617, 619
- Remote Assistance, 531–532
- remote computers
 - managing, 776
 - pinging, 956
- remote connectivity
 - options, 58
- remote controller units, 907
- Remote Desktop Connection, 530–531
- Remote Desktop Protocol (RDP), 530, 531
- Remote Installation Services (RIS), 194
- remote logon, 1046
- removable storage, 233–237
- Removable Storage Management snap-in, 775
- ren command, 622
- Rename command, 595, 931, 932
- repair bench, 13
- Repair command, 931
- repair toolkit, 35
- repairs, 447
- repeat rate, 320
- repeaters, 892–893
- Repetitive Strain Injury (RSI), 44
- required exam, 10
- rescue CDs, 856–860
- reserve resources, 198–199, 370–371
- resistance (ohms), measuring, 468, 469
- resistor in a grounding strap, 39
- resolution, 337, 341, 411–412
- resolution path, 22
- Resource Arbitrator, 691, 692

resource conflicts, 353, 364–365
 resources, reserving, 198–199, 370–371
 restorations, testing, 1031
 restore, performing, 844
 Restore and Manage Media tab, 849
 Restore files, options for, 849
 restore points, 850–853
 Restricted Sites, 995, 996
 retake policy of CompTIA, 17
 revised exams for CompTIA A+, 10
 RFCs (Requests for Comments), 982–983
 RFI. *See* EMI
 RG (Radio Grade) cabling, 879
 RI (Ring Indicator), 307
 ribbon cables, 85, 295, 453
 rights, 1045–1048
 rigid disk number, 649
 RIMM (Rambus Inline Memory Module), 159
 ring topology, 874
 RIS (Remote Installation Services), 194
 riser card in the NLX form factor, 93
 RJ45 port, 79, 80
 RJ-11 connector, 303, 878
 RJ-45 connector, 303, 878
 RJ-45 jack in a laptop, 432
 rmdir (rd) command, 620
 rogue signals, 921
 rollers, laser printer, 381
 ROM (Read-Only Memory), 52, 72, 152–153
 room temperature, 451
 root (.) domain of DNS, 1000
 root directories, 250, 615
 routable protocols, 936, 938

router password for wireless, 1059
 routers, 894–895, 898–899, 943
 RPM (Red Hat Package Manager), 504–505
 RSI (Repetitive Strain Injury), 44
 RS-232 ports. *See* serial ports
 RS-232C standard, 293
 RTS (Ready To Send), 307
 Run registry keys, 754–755

S

Safe Mode, 695, 801, 802, 836–838
 Safe-Mode warning message, 656
 safety precautions for printers, 402
 SAM and SECURITY keys, 842–843
 SAM (Security Accounts Manager), 1025
 Samba Network Neighborhood, 858, 860
 SATA (Serial Advanced Technology Attachment), 51, 85, 86–87, 232–233, 276, 279
 satellite Internet services, 1003–1004
 Save As dialog box, 629
 SC (subscriber connector), 883
 ScanDisk, 261–262, 655
 scandisk.exe, 583, 833
 scandisk.ini, 833
 Scanner and Camera Wizard, 413
 scanners, 407–408, 409, 463
 Scanners and Cameras Wizard, 410
 scanning, 409–411, 749
 scan-only mode, 764
 scanstate, 574
 scatter infrared devices, 390
 schedule(s)
 for preventative maintenance, 448–449
 setting for printers, 394
 for virus scans, 1069
 scheduled definition updates, 1068
 scheduled tasks, running, 782–783
 scope of a problem, 22
 screen flicker, 483
 screen fonts for MS-DOS applications, 720–721
 screen projectors, 340
 screen resolutions for monitors, 57
 screwdrivers, 35–36, 470–471
 SCSI (Small Computer System Interface)
 adapter, 222, 231
 advantages over IDE, 223
 BIOS, 230
 bus, 224, 228–229, 231
 cabling, 227–228
 cards, 229
 connection, 409
 connector, 87, 227
 controller, 87, 649, 796
 described, 51, 346
 devices, 230–232
 drives, 85
 host adapter, 223, 231
 ID numbers, 225–226, 227, 230, 649
 overview of, 222–230
 performance over IDE devices, 223
 printers, 391
 terminator, 228, 229
 types of, 229–230

- SCSI-1, 229
- SDRAM (Synchronous DRAM), 158–159
- SEC (Single Edge Contact) chip, 124–125
- SEC (Single Edge Contact) packaging, 66, 67, 131, 133, 134
- SEC2 chip, 124
- SEC2 package, 132, 133
- secondary controller for EIDE, 215
- secondary hard disk controller, 355
- secondary network, wireless for, 908
- secondary ring in FDDI, 889
- sector blocks, 209
- sector translation, 212
- sectors, 209
- secure shell, 994
- secure tunnels, 994
- securing wireless networks, 914–918
- security
 - best practices, 1057–1060
 - bits, 470–471
 - features not supported in FAT, 543
 - fixes, 1074, 1075
 - fundamentals of, 1013–1032
 - implementing, 1039–1060
 - lack of in the FAT file system, 237
- Security Accounts Manager (SAM), 1025
- security certificates, 992.
See also certificates
- Security log, 777, 778, 806, 1053–1054
- security policy, 693, 1052
- security screws, 471
- security settings, 693, 995–996
- Security+ certification, 5
- seek time, 211
- self-calibration of printing devices, 491–492
- self-replicating worm, 1020, 1072
- Serial ATA. *See* SATA
- serial cables, 307, 388, 900
- serial devices, 346
- serial modem cable, pin configuration, 307
- serial mouse, 76, 322
- serial ports
 - configuring in CMOS, 190
 - described, 59, 292–293
 - disabled in the Device Manager, 485
 - disabling in CMOS, 190
 - on the motherboard, 75–76
 - for a mouse, 322
 - networking two computers, 899–900
 - testing using loopback plugs, 473
 - troubleshooting, 484–485
 - types of, 76
 - verifying, 485
- serial (RS232) cables, 484
- server(s), 867, 993, 994, 1021
- server room, 1020
- Server service, 752–753, 970
- server-based (client-server) networks, 869–871
- server-based networking, 870
- service documentation, 26
- service load process, 652
- Service Pack level, 516–517
- service packs, 799
- service properties dialog box, 752–754
- Service Set Identifier. *See* SSID
- services
 - backing up server, 846
 - causing auto-restarts, 815
 - defined, 970, 992
 - described, 938–939
 - disabling any not required, 750–754
 - within the operating system, 970–973
 - properties for, 752
 - restarting, 970, 971
 - setting to restart on failure, 753
 - during the troubleshooting process, 844
 - turning off by choosing Stop, 751
- Services Administrative Tool, 750–751
- Services MMC snap-in, 812, 813
- Services snap-in, 775
- session hijack, 1018
- session services for TCP/IP, 987
- set command, 635
- setramd.bat, 835
- settings, backing up, 23
- setup boot disks, 538
- setup disks, 368
- setup loader phase of Windows installation, 544
- Setup Manager, 562, 563
- setup program, 542, 823
- setup switches, 545–546
- setup.exe, 823
- setupldr, 544
- Setuplog.txt, 564
- setupmgr.exe, 562
- SFX specifications, 274
- SFX v3.1 (Small Form Factor), 279, 280

- SG (Signal Ground), 307
- SGML (Standard Generalized Markup Language), 990–991
- SGRAM (Synchronous Graphics RAM), 338
- shadow RAM, 156
- shadowing, 156
- shared folders, 961–965, 1046
- Shared Folders snap-in, 775
- shared printers, 396–397
- share-level access control, 959–960
- shares, 959, 961, 962, 966–968
- sheetfeed scanner, 408
- shelf life of media, 461
- shell, 513, 652–653, 662
- Shielded Twisted Pair (STP) cabling, 296, 876, 877
- shift command, 633
- short filename, 581
- shortcut keys, 727
- shortcuts, 714, 821–822
- shorting probe, 42
- Show Processes, 740
- shredder application, 1029
- Signal Ground (SG), 307
- signature notation for a drive, 649
- SIMD (Single Instruction Multiple Data), 131
- SIMM/DIMM sockets, 68–69
- SIMMs (Single Inline Memory Modules), 69, 160–161, 162, 166–167
- Simple File Sharing option, 965
- Simple Mail Transport Protocol (SMTP), 988
- simple volumes, 247
- simplex communication, 898
- single density disks, 478
- Single Edge Contact chip (SEC) chip, 124–125
- Single Instruction Multiple Data (SIMD), 131
- single threaded applications, 738
- single threaded environment, 705
- single-mode fiber (SMF), 883
- single-ported video memory, 338
- slave, IDE drive as, 220–222
- sleep mode, 426
- slipstreaming, 571
- slot covers, 473
- Small Computer System Interface. *See* SCSI
- Small Outline Dual Inline Memory Modules. *See* SODIMMs
- smart battery, 424
- smart card, 1025
- S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology) hard drive, 460
- smartdrv.exe, 669–671
- SMB (server message block) format, 895
- SMF (single-mode fiber), 883
- SMTP (Simple Mail Transport Protocol), 988
- smurf attack, 1017
- snap-ins, 530
- SNMP (Simple Network Management Protocol), 460
- social engineering attacks, 1014–1015
- socket 7 ZIF slot, 91, 92
- socket services, 435
- socket types, 125–127, 138
- socket-based session interface, 987
- socks setting, 997
- SODIMMs described, 162, 163
- inserting, 169, 170, 436
- on laptops, 423
- speed of, 167
- soft power support, 94
- soft power switch, 273, 489
- soft reboots, 478
- soft skills, 19–31
- software, 54
- software components of networking, 927–939
- software diagnostic tools, 474, 803–806
- software modems, 484
- SOFTWARE section of the Registry, 842
- software setup, 368
- software tools, 38
- Software Update Service, (SUS), 1074–1076
- software utilities, 450
- software-based attacks, 1019–1020
- SOHO (Small Office/Home Office) routers, 1005
- Sound and Audio Devices Properties dialog box, 345
- sound capabilities of motherboards, 79, 80
- sound cards configuring, 345
- installing, 344
- I/O address range of, 355
- troubleshooting, 482
- types of, 343–344
- sound device, 345
- sound ports on a motherboard, 79, 80
- sound theme in Windows, 345

- source address for a packet, 1017
- source files, 784
- South Bridge chip, 70, 71
- spanned volume, 247
- S/PDIF (Sony Philips Digital Interface Format)
 - connection, 310, 482
- speaker port, 79, 80
- spectrum analysis tool, 921
- speculative execution, 130
- speech-to-text, disabling, 729
- speed
 - in bps, 485
 - of memory, 166–167
 - of modems, 323
 - of motherboards, 106
 - of Ultra DMA devices, 216
- SpeedFan, 460
- spikes in power, 282
- spin speed, 211
- spinner on a SCSI device, 225–226
- splash screen of Windows, 653–654
- spoofing attack, 1017
- spool directory, moving, 748–749
- spool settings in Windows, 394
- spooler service, 812, 970
- spooling, 394–395, 398, 399–400
- SPP (Standard Parallel Port), 189, 346
- spyware, 1072–1073
- SQL Server, 1077
- SRAM (static RAM), 119, 154–155, 164
- ssh, 994
- SSID (Service Set Identifier)
 - disabling broadcasting, 915, 1060
 - setting, 1059–1060
 - for a wireless network, 910
 - wireless network auditing tools, 921
- SSL (Secure Sockets Layer), 918, 992, 997, 1028
- ST (straight-tip) connector, 883
- stalled print spooler, 812
- standalone applications, 700
- Standard Generalized Markup Language (SGML), 990–991
- Standard Parallel Port (SPP), 189, 346
- standard ports, I/O
 - addresses of, 354
- standard timed exams, 10
- standby UPS, 283
- star topology, 873
- star-bus topology, 874–875
- Start menu, 503, 530
- startup disk. *See* boot disks
- Startup disk, 795
- Startup group, 754–755
- state, 572–574
- static RAM. *See* SRAM
- static shock, 38
- STD (Standards)
 - documents, 982
- stdin (standard input), 632
- stdout (standard output), 632
- Step-by-Step with Confirmation boot option, 840
- StickyKeys, 320
- stop bits, 485
- Stop errors, 817–819
- Stop events, 517
- storage, working with, 207–263
- storage devices, 51–52, 437–438
- storage media, 461
- storage unit, 237
- STP (shielded twisted pair) cable, 296
- straight-tip (ST) connector, 883
- striped volumes, 247, 253–256, 1031
- strong passwords, 1016, 1025, 1026, 1058
- Subdirectory attribute, 595, 596
- subnet mask, 942–943
- subscriber connector (SC), 883
- sulfnbk.exe, 583
- Super VGA (SVGA), 57, 337
- superscalar design, 129
- surge suppressors, 282, 456
- SUS (Software Update Service), 1074–1076
- SUSE Linux, 505
- SVGA (Super VGA), 57, 337
- S-Video, 310
- swap file, 519, 665, 723, 801–802
- swap process, 666
- switch-based keyboard, 316
- switches, 617, 896–897, 1059
- synchronization in WSUS, 1075
- Synchronous DRAM (SDRAM), 158–159
- Synchronous Graphics RAM (SGRAM), 338
- sys command, 627
- sys.com, 833
- sysedit.exe, 673–674

SysInternals Web site, 730
 System attribute, 595, 596
 system BIOS. *See* BIOS
 system boards, 53. *See also* motherboard(s)
 memory synchronized to, 158
 system bottlenecks, 735–736
 system catalog, 602
 System Configuration Editor, 673–674
 System Configuration Utility, 766–767
 System Control Panel, 515
 system events, auditing, 1053
 System File Checker (*sfc.exe*), 781–782
 system files, 592–593, 605–606, 625, 842
 system housing. *See* case(s)
 System Information tool, 515–516, 776–777
 system lockups, 819–820
 System log
 in Event Viewer, 777, 778
 of Windows XP, 806
 system memory, 117, 152
 System Monitor, 741
 System Partition, 647
 system resources
 additional for every NTVDM and WOW, 715
 assigning to devices, 352
 changing, 365–371
 examining, 351–371
 understanding, 352–363
 working with, 363–371
 system restore cache, 851–852
 system restore point, 695, 850

System Restore Wizard, 850–853
 SYSTEM section of the Registry, 842
 system state, backing up, 845, 846
 system time, 1046
 system tray, 530, 754
 system.dat file, 661
 system.ini, 652, 661, 671–672, 800
 systemroot command, 844
 systems
 hardening, 1057
 patching, 1057
 securing through the BIOS, 1039–1040
 systems devices, power usage, 274
 Systray. *See* system tray

T

T1 connections, 1003
 T3 connections, 1003
 TA (Terminal Adapter), 58
 tablet PC, 492
 tape drives, 216, 236–237, 460
 tapes, rotating, 1031
 Task Manager, 714, 736, 737, 819–820, 823
 Task Scheduler, 450
 Task Scheduler service, 782–783, 847
 taskbar, 503, 530
 T-connector, 304
 TCP, 986
 TCP/IP
 configuring in Windows 2000/XP/2003, 944–945
 configuring using DHCP, 945–946
 described, 938, 939–946, 986–987
 planning settings, 954
 telephone conversations over, 891
 troubleshooting connections, 958
 used by the Internet, 982–983
 utilities, 954–958
 TCP/IP configuration gathering, 638
 to network clients, 918
 TCPView, 816
 TD (Transmitted Data), 307, 900
 technician exams, 10
 telnet, 993
 temperature for LCD panels, 428
 temperature sensors, processors with, 142
 Temporal Key Integrity Protocol (TKIP), 917
 temporary files, 750
 Temporary Internet Files, 995
 terminal window, 506
 Terminate and Stay Resident programs. *See* TSRs
 termination of the SCSI bus, 228–229
 terminator, 872, 881
 test page, printing, 395, 398
 testing hypotheses, 20–21
 testing process, 14–17
 text, searching and replacing, 630
 text editors, 772, 794, 798
 text files, modifying, 784
 text-mode phase of Windows installation, 544–545, 557

- TFX v2.2 (Thin Form Factor), 279, 280
- Themes service, 752, 754
- thermal compound, 142
- thermal printers, 387
- thermal sensor, 121, 133, 142
- thermal stress, 451, 456
- thermometers, digital
 - infrared, 471
- thick coax, 297
- thicknet, 878–879, 882, 887
- thin coax, 297
- Thin Film Transistor (TFT) displays, 428
- thinnet, 878–881, 887, 892
- Thinnet networking,
 - connector for, 303
- threads, 121, 709, 738
- throttling, 121
- thumb drives. *See* flash drives
- Thumbnails view, 525
- TIFF file, 410
- time, system, 1046
- timer hardware emulation, 720
- TKIP (Temporal Key Integrity Protocol), 917
- TLS (Transport Layer Security), 992
- ToggleKeys, 321
- token passing, 885–886
- Token Ring, 888–889, 890
- toner, 382, 383, 384, 402
- toner kits, 43
- tools, 35–38, 527–532
- top-level directories, 615
- top-level domains of DNS, 1000
- topology, 871, 886
- Torvalds, Linus, 503
- TORX screwdrivers, 36
- touch pad, 329, 430–431
- touch screens, 329
- trace route (`tracert.exe`) utility, 957
- transfer corona wire, 382, 384
- transfer rates
 - of hard drives, 217
 - of network cards, 326
 - of USB 2.0, 80
- transferring phase of laser printing, 384
- transistors in a processor, 141
- Transmission Control Protocol/Internet Protocol. *See* TCP/IP
- transmissions, securing, 1027–1028
- Transmitted Data (TD), 307, 900
- Travan tape drive standard, 237
- Trojan horses, 1019, 1071
- troubleshooting
 - AC adapters, 427
 - art of, 474–492
 - basics, 475
 - device installation, 695
 - dot matrix printers, 400
 - hardware devices, 683
 - inkjet printers, 401
 - laptop keyboards, 430
 - laser printers, 401
 - learning skills of, 19–27
 - physical environment for, 475
 - printers, 397–401
 - procedures, 19–20, 468–492
 - process, 22–25
 - sections to, 19
 - with TCP/IP utilities, 954
 - tools, 468–474
 - utilities, 759–785
 - wireless networks, 921
- TrueType fonts, 720
- Trusted Sites, 995, 996
- TSRs (Terminate Stay Resident programs), 659, 722, 826
- tweezers, 36, 470
- twisted pair cables, 295–297, 876–878
 - .txt files, 771
- Type A or B USB connector, 304
- type of memory in CMOS, 188

U

- UDP (User Datagram Protocol), 986
- Ultra DMA (Direct Memory Access) drives, 216–217, 218
- Ultra SCSI, 230
- Ultra Wide SCSI, 230
- UMBs, 667–668
- unattend file. *See* answer file
- unattended installations, 561–564
- unattended installation switch, 545, 546
- UNC (universal naming convention) path, 396, 967, 969
- unicast network traffic, 739
- Uniform Resource Locator (URL), 998–1001
- uninstal.exe, 833
- UnInstall key in the Registry, 703
- uninstall routines, 704
- uninterruptible power supplies. *See* UPS
- Universal Data Connector (UDC). *See* IDC connector

- universal naming convention (UNC) path, 396, 967, 969
 - Universal Serial Bus. *See* USB
 - UNIX, 503
 - unplugged devices, viewing, 687
 - unshielded twisted pair cable. *See* UTP cable
 - updating virus definitions, 1070–1071
 - Upgrade Advisor for Windows XP, 567–568
 - upgrade path to Windows 2000, 566
 - upgrading
 - memory, 166
 - Windows OS, 565–568
 - uploading pictures, 413
 - upper memory, 664
 - Upper Memory Blocks (UMBs), 667–668
 - UPS (Uninterruptible Power Supplies), 282–284, 456
 - URL (Uniform Resource Locator), 998–1001
 - U.S. Environmental Protection Agency (EPA), 42
 - USB 1.0, 80, 293
 - USB 1.1, 293, 389
 - USB 2.0, 80, 293, 389, 390
 - USB (Universal Serial Bus)
 - bus, 51
 - chain, 390
 - connectors, 304, 305, 318, 319, 323, 389
 - described, 293, 346
 - devices, 195–196, 235, 486
 - external drives, 235, 437, 438
 - hard drives, 486, 487
 - hub, 81, 195
 - mixing 2.0 and 1.x, 486
 - mouse, 323
 - network adapter, 928
 - printers, 389–390
 - troubleshooting, 485–487
 - USB (Universal Serial Bus) ports
 - connecting scanners, 409
 - described, 59, 80–81
 - disabled, 486
 - disabling in CMOS, 187, 195–196
 - enabled in CMOS, 195
 - IRQ associated with, 360
 - USENET news groups, 984
 - user(s)
 - adding to the Administrators group, 730
 - implementing, 1040–1045
 - limiting connection to shares, 962
 - modifying directory settings, 750
 - setting permissions for, 963, 1026
 - tracking on Web sites with cookies, 996
 - user accounts, 561, 1024, 1040–1043
 - user data files, restoring, 572–574
 - user interface of Windows, 513
 - user mode, 517, 518
 - user password. *See* power-on password
 - user portion of the Registry, 801
 - user rights, 1046–1047
 - user settings in the Registry, 523
 - User State Migration Tool (USMT), 572
 - user32.dll, 661
 - user.dat file, 661
 - user.exe, 661
 - user-level access control, 960
 - Users groups, 1044
 - usetup.exe, 544
 - USMT (User State Migration Tool), 572
 - UTP (Unshielded Twisted Pair) cable, 295–296, 876, 877, 878, 892
 - UUID (Universally Unique Identifier). *See* GUID
-
- ## U
-
- VA (volt-amps), 283
 - vacuums, 449
 - vampire tap, 882
 - vaporized ink, 385
 - vendor training, 26
 - ventilation, 450–454, 458, 476
 - ver command, 627
 - Verify Data after Backup option, 847
 - version numbers
 - of each BIOS, 182, 183
 - for modern versions of Windows, 516
 - VESA local bus (VLB), 100
 - VGA (Video Graphics Array), 57, 337
 - video, 415, 482–484
 - video accelerator, VRAM on, 156
 - video adapters
 - described, 335–339
 - installing additional, 342–343
 - in laptops, 422
 - on the motherboard, 77–78
 - video cards, 52, 57, 339. *See also* video adapters

- video controller. *See* video adapters
 - video driver, 339
 - Video Electronics Standard Association. *See* VESA
 - Video Graphics Array (VGA), 57, 337
 - video memory, 338–339
 - video port, 59, 336
 - Video RAM (VRAM), 156, 338
 - video standards, 336–337
 - views, changing in My Computer, 524–525
 - Virtual Desktop, 531
 - Virtual Local Area Networks (VLANs), 1059
 - Virtual Machine (VM), 664, 707–708, 716
 - virtual memory, 120, 514, 664–666, 746–748
 - Virtual Memory Change button, 519, 520
 - Virtual Memory Manager (VMM), 518–519, 664–666
 - Virtual Private Network (VPN), 918, 1028
 - virtual zoom, 412
 - virus definitions, updating, 1068, 1070–1071, 1076
 - virus hoax, 583, 1073
 - virus protection, 196–197, 826
 - virus scans, 1068–1069
 - viruses
 - altering the MBR, 196
 - defined, 1065
 - described, 1019
 - loading into memory as a TSR, 826
 - protection from, 1067–1068
 - scanning for, 460
 - types of, 1065–1066
 - virus-protection software
 - described, 1067
 - updating virus definitions, 1076
 - using, 1059, 1068–1070
 - Visual Basic Script (VBScript), 590, 783–784
 - visual cues in troubleshooting, 475–476
 - VLANs, 1059
 - VLB (VESA local bus), 100
 - VM (virtual machine), 664, 707–708, 716
 - VMM (Virtual Memory Manager), 518–519, 664–666
 - vmm32 (Virtual Memory Manager), 660
 - vmm32.vxd, 660
 - voice recordings, 345
 - VoIP (Voice over IP), 891
 - volatile memory, RAM as, 153–154
 - volatile RAM chip, CMOS RAM as, 155
 - voltage, 138, 273, 468, 487
 - Voltage Regulator Module (VRM), 122
 - volt-amps (VA), 283
 - volts, 271
 - volume for the sound card, 345
 - Volume Label attribute, 582, 595, 596
 - Volume Label for a partition, 251
 - Volume Shadow Copy, 847
 - volumes
 - creating in Windows 2000/XP/2003, 244–249
 - formatting, 250–252
 - used by dynamic disks, 247
 - VPN (Virtual Private Network), 918, 1028
 - VRAM (Video RAM), 156, 338
 - VRM (Voltage Regulator Module), 122
 - vulnerability assessment software, 1076–1077
-
- ## W
-
- wait time for memory, 165
 - WAN cellular technology, 433, 920
 - WAN (Wide Area Network), 871
 - WAP (Wireless Access Point), 432, 875, 876, 890, 897–898
 - WAP2, 916
 - warm swappable devices, 692
 - watts, 272
 - wear-out period for devices, 475
 - weather
 - affecting wireless, 921
 - satellite Internet susceptible to, 1004
 - Web browsers, 990, 994–998
 - Web cameras, 414, 415
 - Web sites, 1000–1001
 - WEP (Wired Equivalent Privacy), 916, 1060
 - Wide Area Network (WAN), 871
 - Wide SCSI-2, 229
 - Wi-Fi Protected Access (WPA), 916–917
 - wildcards with the `dir` command, 615–616
 - Win16Mutex flag, 710, 711
 - win.com, 653, 660, 793

- Window Accelerator Card RAM. *See* WRAM
- Window Managers, using in Linux, 506
- Windows. *See* Windows OS
- Windows 3.0, 500
- Windows 9x
 - boot disks for, 831–835
 - boot process for, 653–662
 - components used, 661–662
 - disk caching software, 669
 - Dr. Watson with, 803
 - executing in, 707–708, 711–712
 - levels of access control, 959
 - making a bootable disk, 795
 - pressing F8 during boot-up, 836
 - short filenames, 582
- Windows 95
 - boot disk, 831, 835
 - described, 500
 - emulating in Compatibility Mode, 728
- Windows 95 OSR2, 239, 500
- Windows 95 Plus Pack, 833
- Windows 98, 500
 - boot disk, 537, 793
 - bootlog.txt file from, 838–839
 - emulating in Compatibility Mode, 728
 - replacing with Windows XP, 574
 - versions of the boot disk, 834–835
- Windows 98 SE, 500
- Windows 2000, 501
 - 16-bit Windows applications, 712–715
 - architectural diagram of, 708, 709
 - attended installation of, 545–556
 - autoexec.bat file and, 660
 - backup software, 845
 - boot options, 840
 - booting from CD, 547
 - command-line utilities, 613–615
 - configuring for DHCP, 945–946
 - creating a user account in, 1040–1041
 - creating an Emergency Repair Disk, 769
 - default permissions for a share, 962–963
 - Emergency Repair Disk, 841
 - emulating, 729
 - executing in, 708–710
 - extended attributes, 597–603
 - files specific to, 647
 - installing, 547–556, 702–703
 - minimum hardware requirements, 541
 - retrieving the version number, 515
 - setup boot disks for, 538
 - sharing a folder in, 961–964
 - viewing I/O addresses, 354
 - viewing IRQs in use on, 358, 359
- Windows 2000 Professional, 501, 566–567
- Windows 2000 upgrade program, 566
- Windows 2000/XP
 - configuring password complexity, 1058
 - installing a SCSI adapter, 231–232
 - multi-display support, 342
 - viewing memory addresses in use, 363
- Windows 2000/XP/2003
 - bootable CD for, 538
 - built-in groups found in, 1043
 - changing computer names, 947–948
 - configuring as a WINS client, 949
 - configuring multi-display support, 342–343
 - configuring TCP/IP on, 944–945
 - creating partitions, 244
 - creating volumes or dynamic disks, 247–249
 - Disk Defragmenter utility, 260
 - enabling auditing, 1051
 - enabling file and printer sharing, 960–961
 - formatting a partition or volume in, 251
 - installing a network adapter in, 929–930
 - installing a network client, 933
 - installing a network printer, 968–969
 - setup program, 561–563
 - sharing a printer in, 968
- Windows boot disk, 38
- Windows cache size, 670, 671
- Windows Catalog, 541
- Windows CD, 38
- Windows client, configuring, 949
- Windows Defender, 1072–1073
- Windows Device Manager. *See* Device Manager
- windows directory, 647

- Windows Disk Cleanup utility, 750
- Windows environment, optimizing, 735–755
- Windows Error Reporting, 803–805
- Windows Explorer, 513, 525–526, 589
- Windows File Protection interface, 781
- Windows GUI interface, 606–607
- Windows installation, 564–565, 860–861
- Windows Installer, 700
- Windows Internet Naming System server, 949
- Windows keyboard, 317
- Windows ME, 500, 728
- Windows NT, 500, 841
- Windows NT 3.x Workstation, 500
- Windows NT 4.0, 501
 - backup software, 845
 - Enable VGA (Video Graphics Array) mode, 836
 - Plug and Play not supported in, 691
 - similarities with Windows 2000, 566–567
 - supporting the Compress attribute, 598
- Windows NT 4.0 (Service Pack 5), emulating, 729
- Windows on Windows (WOW), 518, 712
- Windows OS
 - 64-bit editions, 134
 - attended installations of, 544–561
 - clients already installed, 932
 - comparing versions of, 500–501
 - default booting installation of, 860–861
 - drive imaging for, 539–540
 - dual booting, 570–571
 - GUI elements, 502–503
 - installing, 537–540, 564–565
 - installing printers in, 392–397
 - looking at, 500–503
 - network installation of, 539
 - in a networking environment, 928
 - preparing for installation, 540–544
 - system requirements, 501–502
 - unattended installations of, 561–564
 - upgrading, 565–568
- Windows PIF Settings dialog box, 719
- Windows RAM (WRAM), 156, 338
- Windows screen saver, 725–726
- Windows Script Host, 783–784
- Windows Server 2003
 - creating a mirrored volume on, 253–256
 - creating a RAID 5 volume, 257–259
 - installing a SCSI adapter in, 231–232
 - multi-display support, 342
 - viewing a list of IRQs in use on, 358, 359
 - viewing I/O addresses, 354
 - viewing memory addresses in use, 363
- Windows services. *See* services
- Windows setup program, 545
- Windows shortcut keys, overriding, 727
- Windows Sockets (Winsock), 987
- Windows Software Update Services (WSUS), 1074–1076
- Windows source files, updating, 571
- Windows Update, 817, 1057
- Windows Update command, 1074, 1075
- Windows Vista, 501
- Windows vulnerabilities, 1077
- Windows XP, 501
 - 16-bit Windows applications, 712–715
 - 16-bit Windows environment, 710
 - accessing the command line, 527
 - applications, uninstalling or removing, 704
 - architecture identical to Windows 2000, 708
 - archive bit, viewing, 1030
 - ASR process, 853–855
 - attended installation of, 556–561
 - autoexec.bat file and, 660
 - backup software, 845
 - boot files and boot errors, 795–799
 - boot options, 840
 - boot process, standard, 646–653
 - booting, files required for, 835

- command-line utilities, 613–615
- compatibility modes, 728–729
- core files required to boot, 796–797
- data sent to `stdout`, 632
- default logs, 806
- digital camera, uploading pictures from, 413
- Disk Cleanup tool, 262–263
- display settings in, 341–342
- DNS server, 951–952
- executing in, 708–710
- extended attributes, 597–603
- files specific to, 647
- firewall feature in, 1055–1056
- folder, sharing, 964–965
- groups, 1044–1045
- hardware requirements, minimum, 541
- installing a new program, 702
- integrity check, 262
- I/O addresses, viewing, 354
- IRQs, viewing, 358, 359
- keyboard options in, 319–321
- LAN connections, viewing, 931
- microphone as a recording source, 416, 417
- MS-DOS-based applications in, 716–728
- NTFS permissions in, 1048
- NWLink, installing, 937
- pressing F8, 836
- printer, installing, 392–394
- resource assignments, 369
- restore points, 850–853
- scanning an image, 409–411
- security policies, 1047–1048
- services part of, 751–752
- short filenames, 581–582
- stability advantage, 712
- stopping, starting, or restarting a service in, 970, 971
- system requirements for, 501
- test page, printing, 398
- upgrading to, 567–568
- user account in, 1040–1041
- version number, retrieving, 515
- Windows Catalog, 541
- wireless configuration with, 910, 911
- Windows Zero Configuration, 910
- `win.ini`, 662, 672–673, 800
- `winmsd` (Windows NT Microsoft Diagnostics), 516
- `WinMSD.exe`, 776
- `winnt32.exe`, 544, 545
- `winnt.exe`, 544–545
- `winnt.sif` file, 563
- WINS (Windows Internet Naming System) server, 948–949
- Winsock interface, 987
- wire cutter/stripper, 472, 473
- Wired Equivalent Privacy (WEP), 916, 1060
- wireless 802.11b, 890
- Wireless Access Point. *See* WAP
- wireless access to the Internet, 1004
- wireless clients, 897
- wireless connection, 910
- wireless devices, 907
- wireless networks
 - auditing tools, 921
 - cards, 432, 908, 909
 - components of, 908–910
 - connecting to, 910, 911
 - defined in 802.11, 890
 - describing, 908–918
 - increasing performance, 920–921
 - securing, 914–918
 - treating as a hostile remote network, 917
 - troubleshooting, 921
 - types of, 890
- wireless routers, 897–898
- wireless, securing, 1059–1060
- wireless standards, 910–914
- wireless topology, 875–876
- Wireless Wide Area Network (WWAN), 920
- workgroup, 541–542, 553, 554–555, 948
- working directory, invalid, 824–825
- Workstation service, 971
- World Wide Web, 983, 990, 1000
- World Wide Web Consortium, 992
- WORM disc, 234
- worms, 921, 1020, 1072
- WPA (Wi-Fi Protected Access), 916–917
- WRAM (Windows RAM), 156, 338
- wrist rest, 45, 317

write caching, 669, 670
Write permission, 1049
writeable CDs, 234
writing phase of laser
 printing, 383
wscript.exe, 590, 784
WSUS (Windows Software
 Update Services),
 1074–1076
WWAN (Wireless Wide Area
 Network), 920

X

xcopy command, 621
Xeon processor, 133–134
XHTML standard, 992
XML, 991–992
XT, IRQs available on, 357
XT keyboard, 317
XWindows, 504

Z

ZIF (zero insertion force)
 socket, 66, 67, 123–124,
 138–141
zip drives, 214, 216, 236
zoom on digital cameras,
 412
Zotob worm, 816