

Weed them out now!
Detect and dispose of rootkits

Rootkits

FOR
DUMMIES[®]

Security
first-aid tools
on the CD-ROM

**A Reference
for the
Rest of Us!**

FREE eTips at dummies.com*

Larry Stevenson
Nancy Altholz



Rootkits
FOR
DUMMIES®

Rootkits FOR DUMMIES®

by **Larry Stevenson and Nancy Altholz**



Wiley Publishing, Inc.

Rootkits 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. 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. FULFILLMENT OF EACH COUPON OFFER IS THE SOLE RESPONSIBILITY OF THE OFFEROR.

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: 2006926390

ISBN: 978-0-471-91710-6

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1

1B/RS/QR/QX/IN



About the Authors

Nancy Altholz (MSCS, MVP): Nancy is a Microsoft Most Valuable Professional in Windows Security. She holds a master's degree in Computer Science and an undergraduate degree in Biology and Medical Technology. She is a Security Expert, Rootkit Expert and Forum Lead, and Wiki Malware Removal Sysop at the CastleCops Security Forum. She has also volunteered at other online security forums. As Wiki Malware Removal Sysop, she oversees and authors many of the procedures that assist site visitors and staff in system disinfection and malware prevention. As a Security Expert and Rootkit Expert, she helps computer users with a variety of Windows computer security issues, including malware removal. Nancy coauthored the *Winternals Defragmentation, Recovery, and Administration Field Guide* for Syngress Publishing which was released in June 2006. She has recently been asked to write the foreword for a book authored by Mingyan Sun and Jianlei Shao, (developers of the DarkSpy Anti-rootkit program), on advanced rootkit detection techniques. She was formerly employed by Medelec: Vickers' Medical and Scientific Division, as a Software Engineer in New Product Development. Nancy's interest in malware and rootkits evolved as a natural extension of her interest in medicine and computers, due to the many parallels between computer infection and human infection. Besides the obvious similarities in naming conventions, both require a lot of detective work to arrive at the correct diagnosis and enact a cure. Nancy enjoys investigating the malware life cycle, and all the factors and techniques that contribute to it – in short, she likes solving the puzzle, and of course, helping people, along the way. Nancy lives with her family in Briarcliff Manor, NY.

Larry Stevenson: Larry has worked as a security consultant for over fifteen years. His education is abundant, including continuing studies in computer security, history, and fine arts. Larry works as an expert, volunteer moderator, and writer on staff at CastleCops, providing assistance and written articles to all users. In 2005, he wrote weekly articles on computer security topics for the *Windows Security Checklist* series. He helped develop, and co-wrote the *CastleCops Malware Removal and Prevention* procedure. For these published efforts he was given the MVP Award: *Microsoft Most Valuable Professional* in Windows Security, 2006. Currently a co-founder with Nancy Altholz of the CastleCops Rootkit Revelations forums, he continues to develop ways for users to obtain assistance and information from rootkit experts. A Canadian citizen, he is currently employed at a multi-function, government-owned facility which includes private residences for people with special needs, a senior citizens care home, daycare center, offices, a cafeteria and a public access theater. For over seven years he has served as the Chief Steward in the union local, negotiating contracts and solving workplace issues.

Dedications

To my mother, Jeanne Gobeo, for being my constant supporter and friend —
and to my sister, Rosie Petersen, for making this world a rosier place. — *NA*

To Lael and Ken Cooper, Tiffany and Kyla, Paul and Robin Laudanski,
also to my Muses, and my parents, Ruth and Hatton, for their faith and
encouragement. — *LS*

Authors' Acknowledgments

We are grateful for the tremendous assistance and unstinting dedication of the many people who contributed to this book, both at Wiley and CastleCops. We would especially like to thank Paul and Robin Laudanski for their extraordinary contributions to computer security in general and the generous ongoing support they extended during the writing of *Rootkits For Dummies*. We give thanks to all the people on the Wiley team for their expertise and patience, including Melody Layne, Rebecca Huehls, Laura Moss, Barry Childs-Helton, James Russell, and Technical Editor Lawrence Abrams (BleepingComputer) for the outstanding job he did. We offer heartfelt gratitude to the Advisors and Rootkit Research Team at CastleCops, every one an expert in their field: Media Advisor Mahesh Satyanarayana (swatkat), Firefox Advisor Abdul-Rahman Elshafei (Abulbrahim), Firewall Advisor Allen C Weil (PCBruiser), IE7 Advisor Bill Bright, and our Rootkit Research Team, including Don Hoover (Hoov), James Burke (Dragan Glas), Anil Kulkarni (wng_z3r0), David Gruno (wawadave), and Michael Sall (mrrockford). We would like to acknowledge Wayne Langlois, Executive Director and Senior Researcher at Diamond CS in Australia, for devoting his time, knowledge, and expertise to the “Tracking a RAT” section in Chapter 9. We’d like to thank Przemyslaw Gmerek, developer of the GMER Anti-rootkit program, for freely sharing his rootkit expertise and allowing us to distribute the GMER Anti-rootkit Program on the *Rootkits For Dummies* CD. We’d like to thank Mingyan Sun, codeveloper (along with Jianlei Shao) of the DarkSpy Anti-rootkit program, for freely sharing his in-depth technical knowledge of rootkit methodology and for giving us permission to distribute the DarkSpy program on the *Rootkits For Dummies* CD.

We would like to recognize and extend a special thanks to Mahesh Satyanarayana for sharing his exceptional technical expertise and so much more, during the development of *Rootkits For Dummies*. Nancy would also like to thank her family and friends for their patience and understanding during the course of writing *Rootkits For Dummies*.

We give special thanks to Forensics Advisor Dave Kleiman (CAS, CCE, CIFI, CISM, CISSP, ISSAP, ISSMP, MCSE), who provided valuable insights to our network and forensics sections, and who also helped get this book up and running by providing much needed hardware. Dave has worked in the Information Technology Security sector since 1990. Currently, he is the owner of SecurityBreachResponse.com, and lead litigation support technician for Secure Discovery Solutions, LLC. As a recognized security expert, and former Florida Certified Law Enforcement Officer, he specializes in litigation support, computer forensic investigations, incident response, and intrusion analysis. He is frequently a speaker at many national security conferences and is a published author of computer books. He is also the Sector Chief for Information Technology at the FBI’s InfraGard and Director of Education at the International Information Systems Forensics Association (IISFA).

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: James H. Russell and
Rebecca Huehls

Senior Acquisitions Editor: Melody Layne

Senior Copy Editor: Barry Childs-Helton

Technical Editor: Lawrence Abrams

Editorial Manager: Jodi Jensen

Media Development Specialists: Angela Denny,
Kate Jenkins, Steven Kudirka, Kit Malone

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: Carl Byers,
Denny Hager, Barbara Moore,
Barry Offringa, Heather Ryan

Proofreader: Christine Sabooni

Indexer: Techbooks

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

Contents at a Glance

<i>Introduction</i>	1
<i>Part I: Getting to the Root of Rootkits</i>	7
Chapter 1: Much Ado about Malware	9
Chapter 2: The Three Rs of Survivable Systems	25
<i>Part II: Resistance Is NOT Futile</i>	35
Chapter 3: Practicing Good Computer Hygiene	37
Chapter 4: Staying Secure Online	61
Chapter 5: Patching and Updating Your System and Software.....	101
Chapter 6: Blurring the Lines of Network Security	117
<i>Part III: Giving Rootkits the Recognition They Deserve</i>	149
Chapter 7: Getting Windows to Lie to You: Discovering How Rootkits Hide.....	151
Chapter 8: Sniffing Out Rootkits	179
Chapter 9: Dealing with a Lying, Cheating Operating System	231
<i>Part IV: Ready for Recovery</i>	301
Chapter 10: Infected! Coping with Collateral Damage	303
Chapter 11: Preparing for the Worst: Erasing the Hard Drive	323
<i>Part V: The Part of Tens</i>	336
Chapter 12: Ten (Plus One) Rootkits and Their Behaviors.....	337
Chapter 13: Ten (Plus Two) Security Sites That Can Help You	347
<i>Appendix: About the CD</i>	355
<i>Index</i>	367

Table of Contents

***Introduction*..... 1**

About This Book.....	1
Things You Should Know	2
What You're Not to Read.....	3
Foolish Assumptions	3
How This Book Is Organized.....	3
Part I: Getting to the Root of Rootkits	4
Part II: Resistance Is NOT Futile	4
Part III: Giving Rootkits the Recognition They Deserve	4
Part IV: Readyng for Recovery.....	5
Part V: The Part of Tens.....	5
Icons Used in This Book.....	5
Where to Go from Here.....	6

***Part I: Getting to the Root of Rootkits*..... 7**

Chapter 1: Much Ado about Malware 9

Some Common Questions (and Answers) about Malware	9
Knowing the Types of Malware	10
Viruses	11
Worms	11
Trojans	11
Dialers	12
Backdoors	12
Spyware (and malicious adware)	13
The Many Aims of Malware.....	16
Rootkits: Understanding the Enemy	19
A Bit of Rootkit Lore.....	19
New Technologies, New Dangers	21
Why do rootkits exist?	22

Chapter 2: The Three Rs of Survivable Systems 25

Formulating Resistance	26
Hackers may not be smarter than you	26
Steps to a Better Security Posture	27
Practicing Recognition	30
Spotting signs of malware	31
Recognizing when the problem isn't malware.....	33
Suspecting that you've been compromised.....	33
Planning for Recovery	33

Part II: Resistance Is NOT Futile* 35*Chapter 3: Practicing Good Computer Hygiene 37**

Before Doing Anything.	37
Using System Restore	38
Backing up your Registry	42
Backing up your stuff with Windows Backup	44
Cleaning Your Windows to Improve Security	46
Everything and the kitchen sink: Loading only what you need at startup	47
Removing unused programs	50
Using the Windows Disk Cleanup Utility	51
Defragmenting your hard drive	53
Using Registry cleaners	57
Controlling Removable Devices	58
Disabling AutoRun.....	58
Turning off AutoPlay on all external drives and devices	59
Scanning boot sectors before using external media.....	60

Chapter 4: Staying Secure Online 61

Good Practices Are a Good Start	61
Choosing your contacts carefully	62
Surfing safely	63
Developing strong passwords.....	69
Establishing limited-access user accounts	70
Using a HOSTS file	72
Bashing Your Browser into Submission	73
Saying no to Java, JavaScript, and ActiveX.....	74
Adding sites to your Trusted zone.....	76

Disable AutoComplete in Internet Explorer77
 Using the New Internet Explorer 777
 Surfing with Firefox instead80
 Staying ahead of the game with SiteAdvisor.....81
 Must-Have Protections Online.....82
 Firewall first83
 Scanners Next95

Chapter 5: Patching and Updating Your System and Software101

Preventing Rootkits by Patching Your Clothes102
 Updating Your Operating System.....103
 Patching, updating, and Service Packing103
 Looking at why you need updates104
 Knowing where you can get them105
 Taking advantage of Automatic Updates.....105
 Guide to Windows Update and Microsoft Update.....106
 Patching and Updating Your Software.....113
 Ways to patch or update your applications.....113
 Watching Internet sources for known
 problems with your applications.....114
 Patching and updating shared computers in heavy use114
 Knowing When You Need a New Computer115

Chapter 6: Blurring the Lines of Network Security117

A Checklist for Improving Security118
 Learning to Love Auditing.....119
 Enabling security auditing120
 Using Windows Access Control.....126
 Editing policies and configuring security.....126
 Making your own security-analysis utility127
 Testing your system against a security template.....127
 Customizing a security template for a network135
 Preventing Attacks by Limiting Access139
 Limiting and controlling physical access140
 Using limited-access user accounts140
 Limiting access on networks141
 Making a business security plan143
 Fooling Rootkits with Virtual Operating Systems144
 Planning Your Defense Against Rootkits145
 Establishing a baseline146
 Preparing Recovery Discs147

Part III: Giving Rootkits the Recognition They Deserve 149

Chapter 7: Getting Windows to Lie to You: Discovering How Rootkits Hide	151
Discovering How Rootkits Hide and Survive.....	151
Keys to the Kingdom: Privileges	153
Knowing the Types of Rootkits	154
User-mode versus kernel-mode rootkits	155
Persistent versus non-persistent rootkits.....	157
Hooking to Hide.....	157
How hooking works.....	158
Knowing the types of hooks.....	159
DLLs and the rootkits that love them	160
Privileged hooks	166
Using Even More Insidious Techniques to Hide Rootkits	171
Direct kernel-object manipulation	171
Trojanized utilities	174
Looking into the Shady Future of Rootkits	175
Hiding processes by doctoring the PspCidTable	175
Hooking the virtual memory manager.....	176
Virtual-machine-based rootkits	177
 Chapter 8: Sniffing Out Rootkits	 179
Watching Your Network for Signs of Rootkits	179
Watching logs for clues.....	180
Defending your ports	183
Catching rootkits phoning home.....	192
Examining the firewall.....	193
Trusting Sniffers and Firewalls to See What Windows Can't	199
How hackers use sniffers.....	200
Using sniffers to catch hackers at their own game.....	200
Testing to see whether your NIC is in promiscuous mode	201
Sniffers you can use	202
Investigating Lockups and Other Odd Behavior.....	206
Accessing Event Viewer.....	206
Making some necessary tweaks to streamline logging.....	207
Inspecting event logs with Windows Event Viewer	210
Upgrading to Event Log Explorer	217
Trying MonitorWare.....	219
Checking Your System Resources.....	222
Matching activity and bandwidth	223
Examining active processes	224
Monitoring CPU cycles	228

Chapter 9: Dealing with a Lying, Cheating Operating System 231

- Rooting Out Rootkits232
 - Cleaning a network.....233
 - Before doing anything234
 - The best overall strategy234
- Scanning Your OS from an External Medium234
 - Microsoft WinPE235
 - Non-Microsoft bootable CDs236
- File-System Comparison from Full Boot to Safe Mode238
- Checkpointing Utilities with Offline Hash Databases240
 - Verifying files with FileAlyzer.....240
 - Verifying file integrity with other utilities243
- Rootkit-Detection Tools.....244
 - Autoruns: Aiding and abetting rootkit detection246
 - Rootkit Revealer247
 - F-Secure BlackLight Beta251
 - IceSword253
 - UnHackMe260
 - Malicious Software Removal Tool261
 - AntiHookExec.....262
 - VICE.....269
 - System Virginty Verifier (SVV).....270
 - Strider GhostBuster273
 - Rootkitty274
 - RAIDE275
 - DarkSpy.....276
 - GMER.....283
- Detecting Keyloggers.....289
 - Types of keyloggers289
 - Detecting keyloggers with IceSword.....290
 - Detecting keyloggers with Process Explorer291
 - Tracking a RAT: Using Port Explorer to trace Netbus 1.60293

Part IV: Readyng for Recovery301

Chapter 10: Infected! Coping with Collateral Damage303

- Deciding What to Do if You’re Infected303
 - Knowing when to give up and start from scratch.....305
 - What happens when the patient can’t be saved307
 - Do you want to track down the rootkit-er, or just recover?.....307
 - Taking measured action.....308
- “My Computer Did What?!”310
 - Saving evidence to reduce your liability310

Preparing for Recovery	318
Cutting off network connection before cleaning out the rootkit.....	319
Planning your first reboot after compromise	320

Chapter 11: Preparing for the Worst: Erasing the Hard Drive 323

Don't Trust System Restore After Rootkit Compromise	323
When a Simple Format and Reinstall Won't Work	325
Erasing Your Hard Drive and Installing the Operating System	327
What you need before you begin this procedure	328
Erasing, partitioning, and formatting	329
Installing Windows XP	331
After you install	333
. . . And beyond	333

***Part V: The Part of Tens*** 336

Chapter 12: Ten (Plus One) Rootkits and Their Behaviors 337

HackerDefender.....	338
NTFSsider	339
Elite Toolbar	339
Apropos Rootkit	340
FU — the Malware That's Also an Insult	341
FUto	342
MyFip	342
eEye BootRoot	343
FanBot.....	343
pe386.....	344
Shadow Walker	345

Chapter 13: Ten (Plus Two) Security Sites That Can Help You 347

Aumha.....	348
Bleeping Computer	348
CastleCops Security Professionals.....	349
Geeks to Go	350
Gladiator Security Forum.....	351
Malware Removal.....	351
Microsoft Newsgroups.....	352
Sysinternals Forum (Sponsor of Rootkit Revealer Forum).....	352
SpywareInfo	352
SpywareWarrior.....	353
Tech Support Guy Forum.....	353
Tom Coyote Security Forum	354

<i>Appendix: About the CD</i>	355
System Requirements	355
Using the CD with Microsoft Windows	356
Installing the DART CD applications	356
How to burn an ISO image to CD	357
What You'll Find on the DART CD	357
Bonus Chapters	358
Anti-malware utilities and scanners	358
Backup and imaging applications	359
System-analysis programs.....	360
Rootkit-detection-and-removal applications	361
Password protectors and generators	362
Downloading tools for compromised hard drives	362
Troubleshooting.....	363
 <i>Index</i>	 367

Introduction

Welcome to *Rootkits For Dummies*, a book written for regular folks who need a better understanding of what rootkits are, what we can do to protect our computers and networks against them, and how to detect and remove them. Like Sergeant Schultz on *Hogan's Heroes*, you may be among those who know “nothing, nothing” at all about them. Even the name *rootkit* may be unfamiliar to you — but soon everyone with a computer and Internet access will know how dangerous these malware programs can be.

First, a bit of myth-busting: Rootkits have a scary reputation — just because they're designed to escape detection by ordinary methods, supposedly they can't be seen or extracted. For most of them, that's balderdash. Rootkits are an extraordinary bit of deviance, to be sure, but they *can* be detected — and removed — using tools developed specifically for those tasks. You may still need the help of an expert, but cleaning out those nasty beasties is possible.

Rootkits For Dummies can help you gain insight into the realm of malware, giving you the knowledge and abilities to assess and develop your own plan to prevent this scourge from ruining your day (or week, or year). Whether you have a standalone computer or have a business network to run as an administrator, this book will show you what you can do about rootkits — and help you secure your system against cyber-criminals and all malware, online and off.

You are about to begin a journey from the basics of malware in general to the complex processes of rootkits. We are your guides, with you every step of the way, as you move toward greater computer security competency. We have done our best to provide the most effective tools available, and we've left markers along the path so you won't get lost. In short, this book is both your passport and roadmap to a new beginning in the never-ending saga of Internet security.

About This Book

In *Rootkits For Dummies*, we offer a handy reference guide. You're not expected to read it from cover to cover — although you're welcome to do so, as it's your book — but rather to open it to the parts that interest you the most and

just start reading from there. The 15 chapters (including two bonus chapters on disc), the appendix, and the accompanying DART-CD (which means *Dummies Anti-Rootkit Toolkit*, a CD of tools and utilities to help you protect and clean your computer) provide all the topics and tools essential to dealing with rootkits and their payloads. We wrote each chapter so it could be read on its own; feel free to open the book anywhere and start reading.

Things You Should Know

Although this book comes with a glossary so you can look up what a lot of stuff means, we have some special terms and items we'd like to point out for you just in case there's any confusion or controversy over what things mean in the contexts where we use them.

- ✔ **Blackhats, whitehats, and some maybe gray:** In the old Western movies, the bad guys wore black hats and the good guys wore white ones; it's the same thing here. When we call something black in this book, we usually mean it's bad (if it isn't, we'll tell you); white is good, and gray is slimy.
- ✔ **Hackers and geeks:** These guys are not all created equal. Nothing is wrong with being one, it just depends on what's done with the knowledge of how to hack. We mean no disparagement of these many fine individuals who are good people with brains and skills; if we occasionally use the term "hacker" to refer to a blackhat hacker (see the next bullet), don't hate us. In the old days, to be a *hacker* was a matter of pride and accomplishment. Rather than get involved in these old issues, we decided to be upfront about it from the start. We consider ourselves whitehat hackers, too, and we know they exist and help protect us from the blackhats.
- ✔ **Blackhat hackers:** We consider these to be cyber-criminal hackers, people who use hacker tech and skills for evil purposes, compromising and hijacking people's computers and invading networks with malware and rootkits. These creeps give regular hackers and whitehat hackers a bad name.

Black hat conferences: These shindigs are now held every year (since 1997) at various locations around the globe — featuring cutting-edge security research provided by top business professionals, government security experts, and members of the anonymous hacking communities. These are *good guys*, not a bunch of blackhat hackers! Learn more at the following URL:

www.blackhat.com/main.html

What You're Not to Read

Not that we'd dictate that. It's just that we know your time is precious. To get the essential goods on rootkits and the malware they lug around with them, you don't have to read every single word in this book. Understanding rootkits does take some time, so go ahead and flip through the book. Sidebars and special-information items are provided to help you, but may not be essential to your overall understanding of rootkits — or they may simply be over-the-top technical (you'll know those when you see the Technical Stuff icon). If you're a beginner, or have no immediate interest in this extra material, skip it. (Of course, many techies reading this book will be delighted by these tidbits — and to them we say, *bon appetit.*)

Foolish Assumptions

Most everyone has heard that line about pleasing (or fooling) all of the people all of the time. Well, we aim to please — no fooling — but we also had to make a few practical assumptions about our readers when we started this book. We assumed that you

- ✔ Are familiar with using Windows computers.
- ✔ Know why you need a firewall and antivirus software.
- ✔ Have encountered some form of malware at some point in your adventures with computers, or at least have heard of someone who has.
- ✔ Are getting worried about Internet security on your personal computer or network.

How This Book Is Organized

We have arranged the chapters in this book in five parts. Each part focuses on a particular area of concern to you, the computer user, when you're dealing with malware and rootkits. The book is set up to be eclectic; no need to plow through it in a linear, plodding-along fashion. Play hopscotch with the parts, if you choose: this book was written as a reference, not as a textbook. That said, there is a logical order to the book's parts and chapters; prevention is discussed early on; the identification of rootkits and dealing with the havoc of an infected system are topics introduced later. If you want a full overview, feel free to go the cover-to-cover route.

Part I: Getting to the Root of Rootkits

The book starts by introducing you to malware, rootkits, and the issues they create: what you can expect from rootkits and malware, where you will find it lurking on your system or network, and why you need to know these things. Most networks and standalone computers are ill-equipped to handle the fullest implications of malware and blackhat hacking today. So this part makes no bones about the bad news; you'll discover the plethora of opportunities that cyber-criminals have at their whim, with little or nothing to deter them. Laws have geographical boundaries — unfortunately the Internet does not.

This part provides an overview of the many attacks and malware being encountered on the Internet every day. Before you can secure your computer or network, you need to know what you're up against — malware and rootkits — and the cyber-criminals who use them.

Part II: Resistance Is NOT Futile

This part details the challenges of shoring up your defenses and hardening your computer and network security. From cleaning up the junk languishing in the dark recesses of your computer's file system to using anti-malware applications, you get a handle on what all the geeks and techies already know: By maintaining a clean, balanced, and hardened computer, you can save yourself a lot of hassle, both electronic and financial.

For those who have often felt mystified about how to set up security policies — using either the Local Security Policy Editor (for standalone Windows XP Professional computers) or the Security Configuration Manager (for global network policies), this part is for you.

Part III: Giving Rootkits the Recognition They Deserve

... which is to say, efficient detection, speedy removal, and savvy defense. For both standalone and networked computers, this part shows you how to detect, determine, and remove rootkits. For those of you who like to cut to the chase, here you find the meat of the matter — and an edge you can apply to it (we can already hear you groaning out there!): Here we reveal how rootkits do their special dance, how you can discover them, and how you can put a stop to them.

Part IV: *Readying for Recovery*

Rootkits are nobody's harmless prank; they're often used by cyber-criminals seeking nefarious financial gain. Due to their nature, rootkits can make it difficult to trace the blackhat hacker who put them there. And if they entangle your computer or network as part of a criminal enterprise, you've got potential big trouble. So this part details your options if a rootkit has taken up residence — and shows you what to do about it once you decide on a course of action.

Okay, it had to happen sooner or later: Some rootkits and their malware payloads can so thoroughly compromise a computer that (short of a direct missile strike) they're impossible to remove by conventional means. Even now, many security people claim that you need only reformat your hard drive and reinstall your operating system to get rid of rootkits. Unfortunately, that doesn't work if you have rootkits squatting in the bad sectors of your hard drive. So this part shows you how you really *can* remove even those tough nuts — no missile required — and start over with a clean hard drive.

Part V: *The Part of Tens*

Every *For Dummies* book has a Part of Tens, and this one is no exception. In this part, you get a look at some of the most current rootkits (and a few tough old customers, too), ways that you can protect your computers and networks from them, and the best and the brightest security Web sites that can help you at no charge.

After The Part of Tens you find the Appendix, which gives you an overview of the software available to you on the included CD (as are two more chapters).

Icons Used in This Book

The following paragraphs (with their representative icons) give you an idea of what to expect when you see these icons in the book.



Like torches guiding your path, these icons illuminate special areas for your attention, increasing your wisdom or just making the path a little easier.



Both a heads-up and an FYI, this icon can help guide you on your journey by reminding you of important tidbits to keep in mind.



Danger! Thin ice! Proceed with extra caution when you see this icon. It means what it says. Some procedures are not undoable — especially in this book, where horrors such as reformatting your hard drive are discussed often — but they do require extra care. Slow down and take your time.



Well, yeah, rootkits *really are like rocket science*, extremely technical — but we've done our best to get you up to speed without parboiling your brain. Even so, we feel that some technical details are worth mentioning. If you want a peek under the hood, here you go, but rest assured: You don't have to read this particular stuff.



Whenever an application is featured on the DART CD that comes with this book, you'll see this icon.

Where to Go from Here

One of our favorite ways to see if a book is any good is to open it anywhere but to the first chapter and start reading. If it holds you for more than a page or two, then you know the book is worth your time. Try it yourself here with this book. You can look at the Contents at a Glance page or at the Table of Contents and see what catches your eye, flip to an interesting section, and give it a read. We're proud of the book, and we bet you'll like what you see. So start flipping pages and enjoy this journey as you discover rootkits and how to protect yourself from them.

Part I

Getting to the Root of Rootkits

The 5th Wave

By Rich Tennant



"This rootkit is a little nastier than I expected."

In this part . . .

Cyberspace is a battleground, where computers and networks are invaded, lost, or saved every day. Grab your virtual helmets and gear and let's go take a look at the enemy. You need to know about them, what they do, and why they do it. Until recently the struggle had been more or less equal, but now the enemy has a new and more powerful weapon — the rootkit.

Rootkits keep everything out of sight, invading computers from behind their own lines, acting as delivery systems for the other weapons the enemy uses. Combat with them can be difficult, but not impossible. Learn from battle-seasoned veterans how to survive and win in the war against malware.

Chapter 1

Much Ado about Malware

In This Chapter

- ▶ Posing and answering common questions about malware
 - ▶ Understanding the types of malware (the enemy)
 - ▶ Figuring out what the malware is after
 - ▶ Discovering what rootkits do and why they exist
-

Rootkits have their origin in the Unix world. They were created to replace standard Unix tools with versions that gave a user *root* or super-user privileges, while allowing their activity to remain invisible to other users. A rootkit's unique hiding ability was quickly seized upon by hackers with ill intent as an ideal way to provide cover for devious activities.

If you find a rootkit on your computer, you can pretty much be assured that something else is lurking there, but you won't know what that something is. As malware, rootkits are considered to be among the most insidious and pernicious programs because of their ability to conceal the unknown.

In order to secure your system from rootkits, you need to understand the fundamentals of malware. In this chapter, our goal is to fill you in on those truths, and clue you in to the different types of malware and its aims, as well as the basics of rootkits.

Some Common Questions (and Answers) about Malware

A few questions are quite common when people first hear about malware and rootkits; this section lists the main questions — and, more importantly, the answers to them.

- ✔ **What is malware?** The term *malware* is short for *malicious software*. Malware is created with the intent to enter, modify, or damage the other software on your computer without your knowledge and consent. Like other malware, well-crafted rootkits do all these things — yet remain entirely invisible to the computer user.
- ✔ **What's the relationship between rootkits and malware?** Rootkits' relationship to malware is twofold: To put a rootkit on a computer, other malware has to load it. And after the rootkit is loaded, it's often used to hide more malware. Rootkits created with malicious intent (some rootkits are benign or even beneficial) collectively make up a specific category of malware; however, not all malware programs are rootkits.
- ✔ **Who's vulnerable to malware?** Any computer or network connected to the Internet is a viable target for a malware or rootkit attack. If you are on a broadband or T1 connection, which allows for rapid transfer of data, then you become an even more attractive target to blackhat hackers (about whom more in a minute). Public computers are also vulnerable; someone could just walk by, slip in a disc, and install malware that way.
- ✔ **Who's responsible for malware and what do they want from me?** Malware programmers are often portrayed in the popular press as malcontents, angry at the world, expressing their frustrations with destructive behavior and activities. Although this can be true, the people behind malware are more likely trying to manipulate millions of people, governments, even the stock markets — ultimately in order to make money. The worst among them are criminal and terrorist organizations who exploit the often-lighter sentences imposed for Internet offenses to make pots of money — using malware to steal identities, put the squeeze on Internet-based companies with Distributed Denial of Service attacks, and disrupt commerce with other costly exploits. See the “The Many Aims of Malware” section later in this chapter for more information about what, specifically, those who write and spread malware want from you and your computer.

Knowing the Types of Malware

When you go up against malware, you need to know your enemy. In the sections that follow, you find out about the different types of malware that you need to protect your system from.



Rootkits can be used with any of the major forms of malware described in the following sections.

Viruses

A *virus* is a small program that inserts itself into other executable software. Every time that software is opened and used, the virus program will run, making copies of itself to insert into every document and executable file opened. This can cause damage to your computer software, including your operating system, by corrupting existing data on all your storage media and overwriting your files.

As long as a virus program is present in any software you open, it can spread to other computers when you share files and programs with others — over the Internet using e-mail or P2P (peer-to-peer) file-sharing networks, or via infected CDs, DVDs, or floppy disks. Viruses persist primarily in stored memory on physical media such as your hard drive. New viruses are not as common a threat now as in the past, but they can have the rootkit technology included in their designs.

Worms

Worms are programs that can copy themselves; they exist in RAM (random-access memory). They spread by sending themselves via e-mail, instant-message programs, and peer-to-peer (P2P) file-sharing networks to other computers in a network. Unlike viruses, worms do not insert themselves into other programs — and they rarely affect the files on your hard drive. Worms cripple computers by congesting the flow of information, slowing down the system by using up its resources, or crashing the system altogether — all by making multiple copies of themselves. *Unpatched* computers, — those without the software fixes that plug security holes, — are a bonanza for them. Worms have shut down large portions of the Internet, causing millions of dollars in damages before they were stopped. They can also be carriers of root-kits, backdoors, and trojans (which we describe next).

Trojans

Trojan Horse programs (now mostly referred to as just *trojans*) are malicious applications masquerading as something helpful or innocuous. Veritable “wolves in sheep’s clothing,” they can disguise a destructive program as something more benign, such as an image file. A harmless-looking `.gif` extension, for example, may hide the `.exe` extension of an executable file.

This treacherous type of program was originally called a “Trojan Horse” after the giant “gift” horse (with soldiers inside) that the ancient Greeks offered as a ploy to get inside the city of Troy in *The Odyssey*. In this case, the “soldiers” are executable files — invading programs.





Beware of program files with double filename extensions. By default, Windows hides double extensions.



To make sure you can see double extensions in Windows XP, you need to change just one setting. Here's how:

1. **Click Start, Control Panel, Folder Options.**
2. **Click the “View” tab, and then click Hidden Files and Folders ⇔ Show Hidden Files and Folders.**
3. **To see filename extensions, uncheck the box beside *Hide file extensions for known file types*.**
4. **Click “Apply” and then click OK.**

Now Windows will show all the extensions associated with each file.

Trojans can be contained in a Web site link if you haven't set your Web browser to block scripts. They can also come in as e-mail attachments that you open without scanning first, or be bundled with a program you download from the Internet. Whichever way they reach you, they usually require some action on your part to be installed on your computer.

Dialers

Two kinds of *dialers* exist — one good, one bad. The good one is installed as part of your operating system; it helps you connect to the Internet via an analog dialup connection. The other is malware, used to set up a fraudulent connection (usually to an expensive, long-distance telephone number) or to force downloads — all of which gets charged to your telephone bill — through particular Web sites. Malware dialers can be installed by trojans, ActiveX and JavaScript scripts, and from opening attachments in spam e-mails. (Users of DSL or Broadband connections are usually not affected by dialers.)

Backdoors

Backdoors are programs (or modifications to existing programs) that give outside users remote access to your computer without requiring user identification. Backdoors attempt to remain hidden or to “hide in plain sight” by appearing to be innocent. They can also be special passwords set up on a login system to the same effect.

Backdoors can be installed through weaknesses in an unpatched or unprotected Windows computer, either directly by blackhat hackers or with a trojan, virus, or worm. They can even be installed as “Easter eggs” by the original programmers of software (a practice considered highly unethical).



Easter eggs are hidden programs within software that can be triggered using specific commands. Professional programmers tuck them inside commercial software and then tell other programmers how to access them to get amusing animations or messages. But that “little something extra” can just as easily be malware.

Spyware (and malicious adware)

Currently considered to be one of the greatest threats to Internet and computer security today, *spyware* includes a wide range of applications that use stealth and trickery to fool users into installing them. Broadly speaking, *spyware* takes full or partial control of computer operations while denying your rights to privacy and to choose for yourself what runs on your computer — all for the benefit of strangers. Whether used “legitimately” or illegally, *spyware* is a way for malicious people to attempt to control, monitor, and profit from you against your wishes. (We discuss the aims of malware in more detail a little later in this chapter.)

Adware programs are often associated with *spyware*, because many *adware* programs monitor your browsing habits to target you with specific advertisements. The companies that provide these often-surreptitiously-installed bits of software are quick to point out that their programs are “not *spyware*,” but it’s really six of one, or half a dozen of the other. Legitimate *adware* programs differ from illegitimate applications; they only include advertisements as a way to offset their production and maintenance costs. Illegitimate *adware* bombards you with flashy pop-up ads that won’t go away till you click a Close button (which may trigger more).

Some *adware* programs disguise themselves as beneficial toolbars or search aids when they are anything but that. Such *adware/spyware* tool bars can redirect your browser, bias your search results, or serve targeted pop-up advertisements. There are, of course, toolbars that are legitimate such as the Google Toolbar which do deliver on their stated promise. As a general rule of thumb, legitimate toolbars are easily removable through the Add/Remove programs feature of the Windows Control Panel. *Adware* toolbars are often a nightmare to remove, and often appear out of nowhere on your desktop.

The CastleCops Security Forum maintains a Toolbar research database, which can help you decide whether a toolbar is legitimate or not:

www.castlecops.com/CLSID.html



Just so you know: Legitimate applications are not spying on you, not reporting back to their companies, and not wasting your time by requiring you to close ad windows. By contrast, many illegitimate adware programs provide targeted pop-up ads and build marketing profiles on each user — without the user's knowledge or consent — that can then be sold to other advertising agencies.



You may also know of spyware applications that are considered legitimate and are commercially available. Typically these are for use in specialized circumstances, such as when a company secretly monitors the activities of its employees; parents do likewise with their children who use the family computer, schools monitor their students while online, and so on. Check the laws in your area before using such applications yourself. One of the authors know people who have permanently ruined their relationships with family, friends, and neighbors by using spyware on their computers to monitor their children (this is different from a parental control program). When you spy on your children, you are also spying on their friends. Spying on someone over whom you have no authority is also a crime in most jurisdictions. Employers and institutions can do it, but individuals or parents should avoid these applications entirely. They are like a Pandora's Box. Curiosity can kill your reputation.

Spyware is generally installed in the following ways:

- ✓ **Presenting the spyware as something it's not:** Usually these types of spyware and malicious adware are packaged in a way that offers a perceived benefit to you, such as
 - Helping you search the Internet for Web sites you want to view
 - Providing you with a special program that promises to increase download speeds
 - Pretending to remove a nonexistent spyware threat while creating a real one
- ✓ **Tricking you into believing that a user action is required:** This devious approach may provide (for example) a link that says `Click here to have all media content displayed on this page` — and after it's too late, you realize that your click enabled the installation of an unwanted program.
- ✓ **Bundling the spyware (something you don't want) with a program you do want:** Unlike the preceding example, you do in fact get the program you think you're getting — but you *also* get spyware programs you

didn't necessarily bargain for. Often, people actually agree to download these programs by accepting the program's license agreement. If you actually *read* the entire agreement (which few people do), you may find some legalese that mentions that by downloading this program, you *also* agree to download other programs bundled with the software. The agreement may not tell you what those "other programs" do — but (unfortunately) they may very well be spyware.

- ✔ **Peer-to-peer (P2P) file-sharing programs are a major vector for bundled spyware.** Although not all P2P programs come with a spyware payload, many unfortunately do. Furthermore, the practice of opening your computer to anonymous downloads can introduce additional malware to your computer from infected shared P2P folders. You have to ask yourself whether free is really free, and if the risk of acquiring a rootkit or trojan is really worth the trade off. Early versions of Kazaa, for example, included spyware.



A freeware program called EULalyzer scans the *end user license agreement* (or EULA) of a program for "interesting words and phrases" that might need a closer look. It does not dispense any legal advice, but it helps translate convoluted terms that can crop up in long EULAs. You can download it at

www.javacoolsoftware.com/eulalyzer.html

- ✔ **Installing a connection that automatically downloads additional crud.** The connection is totally dependent on the provider of the malware, and is typically achieved by installing a backdoor (for a rootkit), or a Browser Helper Object (BHO) for ordinary spyware, though some overlap may occur. The connection is then used to download additional unwanted software or updates to existing software to further compromise the infected machine. Usually these remote transfers run in the background, and may only catch your attention by slowing down your Internet access and your computer.
- ✔ **Doing "drive-by" downloads:** In effect, this technique (also known as a *WMF (Windows Metafile) exploit* denies users the right to choose what to put on their computers by installing something they *didn't* choose. A *metafile* contains a bunch of instructions for what and how to display a graphic image. A *drive-by download* is accomplished when you browse to a malicious Web site that uses vulnerabilities in your browser and operating system to force the spyware onto your computer.



A too-easy way to get a drive-by download is to be online without a firewall. You can even get one from legitimate sites that have been hacked to provide malware-based advertising (or their ad-servers might pass along the drive-by in ignorance). By far the most common drive-bys occur to people who either cruise pornographic sites for thrills or fall for scams that send them to *spoofed* (carefully faked) Web sites. Bottom

line: The dark side of the Internet is just as dark as a big city downtown at night; getting a drive-by is like being mugged. The download uses vulnerabilities in unpatched operating systems and browsers — which is another good reason to get Microsoft updates. In addition to Internet Explorer, other kinds of browsers (such as Mozilla's Firefox or Seamonkey) need regular updates for the same reasons.

The Many Aims of Malware

In the past, the majority of computer hackers used to be content to create mischief and leave a signature of their work as a memento of a successful break-in. The more ruthless ones might destroy data or your operating-system files, or even corrupt your BIOS (the computer's setup information), making a reformat and reinstall inevitable. Their primary reward for such activities was essentially the challenge and conquest. They did it because they could.

The seedier aspects of the cyber-landscape have changed considerably in recent years. Malware thrill-seekers still exist, but today, most purveyors of malware are in it for financial gain. Anything that enables them to make money is fair game. Many operate far enough outside the realm of legitimacy to qualify as cyber-criminals. Rootkits in particular are a perfect tool to use in these exploits, because rootkits allow long-term continued access to your computer without detection.

The goals of malware are many — none of them good for you, the user. In the following list we describe the different goals of malware.

So what are these malware coders after? The answer may include any of the following:

- ✔ **Data about your Web surfing:** By tracking your Web habits, they know what your interests are and what advertising should appeal to you in light of your browsing habits. Such spying enables commercial adware companies to serve targeted pop-ups suited to your personal preferences.
- ✔ **Control over your Web surfing:** In an even more invasive twist, your browser Start and Search pages may be hijacked to a Web site of the malware writer's choosing. If your browser is hijacked, then whenever you attempt to surf the Web, you're redirected to a Web site that bombards you with pop-up ads that the unscrupulous affiliate advertisers hope you'll click. Sometimes your browser remains frozen at a Web site where you will become a captive audience for an advertising campaign. When this happens, your entire surfing experience becomes defined by the adware infection.

A bunch of strangers you'll never know nor meet will benefit enormously from your new enslavement. They get their money from the agencies hired by companies to promote their products and services with advertising. No matter how the advertising is promoted (or how sleazy a technique this is), a certain percentage of the entrapped users *will* buy — increasing sales — always.

- ✔ **Your sensitive personal information:** Blackhat hackers may want your personal details to commit identity theft, enable bank-account access, or put fraudulent charges on your credit cards. Among the many ways they might try to get your information are the following:



- **Deciphering weak passwords:** A weak password will allow an intruder easy access to your computer or network. This literally opens the door to all sorts of malicious activity and (in the case of a network) essentially guarantees access to many more computers. That's why using a safe-password generator and protection system is so critical. (We include such programs on this book's CD.) Flip to Chapter 4 for a refresher on how to make stronger passwords, and see the Appendix and Bonus Chapter 2 for more information on the password-related applications we have included on the CD.

- **Using false security alerts to goad you into purchasing a program with hidden malware:** Some trojans may try to scare you by claiming that your computer is infected, when your computer is actually infected *by* the trojan they just planted on it!

Your natural inclination will be to click the warning “bubble” — but *don't*. That click directs you to a bogus antispyware or antivirus Web site — which then attempts to con you into purchasing a useless “security” program to “remove” the nuisance threat. To make this scheme even more convincing, the security alerts intentionally mimic those of Windows, so victims are often fooled into thinking that the real Windows Security Center (instead of a cyber-swindler) is posting the alert. Deception and audacity reached a peak when the Vundo trojan used a near-perfect pop-up fake of the Windows Online Safety Center to redirect users to the Web site for the rogue WinFixer program. (Guess what it didn't fix.) The original WinFixer program is now known as WinAntiSpyware 2006 or WinAntivirus Pro. Same purpose, different name — and twins, no less.

Here's an online article with more information about schemes that try to annoy users into parting with their money in exchange for junk software:

www.websense.com/securitylabs/docs/WebsenseSecurityLabs20052H_Report.pdf



- ✔ **Using your system as a cloak for scam operations:** Some blackhat hackers want to hide behind your system and secretly put your computer or network to work for them. This is done by opening and maintaining an

Internet connection between your system (the server) and remote client computers controlled by the bad guys. *Remote-access trojans* (RATs) are used to commandeer your computer from the remote client by maintaining connections with an open, hidden port they have created. Once a RAT sets up shop, your system can be used for any number of nefarious tasks. In addition to identity theft, black marketeers can use your computer for anything — perhaps as a drop for illegal images or as a zombie for Distributed Denial of Service attacks against the Web sites of other businesses. A *zombie* is a computer slaved to an invisible network that attacks Web sites. When thousands of zombies are used in an attack, it's called a *Distributed Denial of Service* (or *DDoS*).

Cases of malware installed by individuals acting alone do exist, but the greater threat to your life and liberty come in from (believe it or not) the cyber-version of the black market — and its sleazy cousin, the gray market:

- ✔ **Black-market groups are usually underwritten by criminal organizations who will go to any length to achieve their goals.** This includes using malware to record and transmit your personal information and financial transactions, and acquiring your passwords and debit- and credit card numbers. They know how to take you to the cleaners and then some. For example, with the right information, they can take out loans in your name, run your credit cards up in the twinkling of an eye, and clean out all your bank accounts.
- ✔ **Gray-market groups operate specifically to make money by using adware and spyware to promote advertising.** Some call this crew “corpirates,” which succinctly describes what these people do. They can operate as regular businesses or corporations because their methods are less dramatic (and technically more legal) than those of the black-market groups. Secrecy and deception, however, are important parts of their work. Many of these groups provide fake security applications to the public — which then don't perform as expected, but deliver targeted pop-up advertisements to your computer instead. Once installed, such software is often hard to remove — and its Terms of Use are as convoluted as they are compromising to the rights of the computer user.

Many Internet businesses are mostly unregulated, unlike offline ones. Even though they are supposed to adhere to the laws of their countries of registration, they do pretty much whatever they like. Unsuspecting users who expect to be dealt with fairly online . . . are under a false impression. On the Internet, as in the old Wild West, (almost) anything goes! To learn more about these modern cyber-corpirates, please visit the SpywareWarrior Security Web site at

www.spywarewarrior.com/rogue_anti-spyware.htm



The Wild West aspect of online life even shows up in the common terms *blackhat* for malicious programs (and programmers) and *whitehat* for legitimate ones — reminiscent of the headgear worn by (respectively) bad guys and good guys in old Western movies.

Rootkits: Understanding the Enemy

A *rootkit* is a program designed to hide not only itself, but another program and all its associated resources (processes, files, folders, Registry keys, ports, and drivers). Rootkits can be *whitehat* (well-intentioned in purpose but still a potential security risk) or *blackhat* (malicious in nature). Malicious rootkits are often used to compromise and maintain remote control over a computer or network for illegitimate, — often criminal — purposes. Malicious rootkits do their work by hiding malware that installs a backdoor to allow an attacker to have unlimited and prolonged access to the infected computer.

A rootkit infection introduces a fundamental flaw into computer systems: Suddenly you can't really trust the integrity of the operating system or have any faith in the results it reports. Because of this flaw, you may be unable to distinguish whether your systems are pest-free or harboring some uninvited "visitor" that traditional scanners are unequipped to deal with.

When you go up against rootkits, you need to know your enemy. This section gives you the skinny on why they hide, how they survive, and why the little creeps exist in the first place. Chapter 7 discusses the more technical side of rootkits, describing in detail how they hide.

A Bit of Rootkit Lore

Rootkit technology is not new. In fact, rootkits have actually been in existence for over a decade. They were first developed for use on Unix-like operating systems (Solaris and Linux), and later evolved to encompass Windows platforms as well. The first public rootkit developed for the Windows NT platforms made its debut in 1999 when it was introduced by Greg Hogg, a well-known security researcher and owner of rootkit.com. The unusual moniker *rootkit* is actually derived from *root* — a Unix reference (which implies root-level access to a system and administrator privileges) — and *kit* (which refers to the collective set of tools used to obtain that hidden and privileged access).

The discovery of the Sony Digital Rights Management (DRM) Rootkit by Mark Russonovich of Sysinternals suddenly thrust rootkits from relative obscurity to a position of prominence. Until the recent publicity barrage, rootkits had commanded little attention and had been implicated with a relatively small percentage of malware infestations. They were considered an intriguing but rarely encountered curiosity than an imminent threat. Enter the Sony rootkit exposé on October 31, 2005 — and suddenly rootkits took center stage. The Sony rootkit controversy has not only heightened public awareness, but it has also spurred the development of new rootkit technology and research, as well. These days, rootkits are regarded as a real and growing potential threat — and the security community has responded to this upgraded threat accordingly.

This unfolding scenario was bound to happen. As security vendors provided increasingly better solutions to combat nearly every type of pest, malware writers have responded by creating a stealthier and more tenacious breed of malware. Your basic Catch 22 scenario has developed. These new exploits are designed to outfox today's highly refined malware detection and removal programs. By embracing rootkits and their stealthy capabilities, cyber-criminals have found a “new and improved” way to launch an attack.



Stealth programs and rootkits represent a looming threat and the tide of the future. In fact, eweek's December 6, 2005 issue has reported that “More than 20 percent of all malware removed from Windows XP SP2 (Service Pack 2) systems are stealth rootkits, according to a senior official in Microsoft Corp.'s security unit.” A more recent paper by the Microsoft Anti-malware Team entitled “Windows Malicious Software Removal Tool: Progress Made, Trends Observed” published on 6/12/2006, reports a more modest rootkit incidence of 14 percent. When the Sony DRM WinNT/F4IRootkit is factored out the figure drops to only 8 percent. Before you jump with joy over the apparent decrease in rootkit prevalence, let's put this in perspective. The June 2006 statistics represent incident rates on Windows 2000, Windows XP, and Windows Server 2003 computers, as opposed to only the extremely popular Windows XP SP2 platform. This would tend to lower the 2006 figures. The December 2005 statistics are not adjusted to exclude the Sony DRM rootkit and were released soon after its public discovery. It is likely more computers were affected by the Sony DRM rootkit at that time, and that would inflate the 2005 figures.

Microsoft has taken this threat very seriously. Apart from its rootkit tools (currently in development), it has incorporated rootkit detection and removal into a handy program called the Malicious Software Removal Tool (MSRT). A newly updated MSRT is delivered along with Windows updates every month — and it silently scans in the background for several commonly encountered rootkits. (Trying to root them out, so to speak.) In addition to rootkits, the MSRT also scans for some of the most pernicious but prevalent backdoor trojans and worms that known to be out there.

New Technologies, New Dangers

If you're like most of us, you may have faced many of the threats out there in cyberspace, putting security measures in place to protect your system from intrusion (and to remove any malware that does find a way to get in). It's true that many tools perform this function quite successfully when used in combination. But the fact that you're reading this book indicates that you may not be content with those security measures — or even confident that they're protecting you. If that's the case, you're right to be concerned.

With the appearance of rootkits on the scene, none of the brilliant tools developed for recognizing and removing malware threats can perform this function accurately. A rootkit can blind traditional security tools to the presence of malware programs, letting the invaders function unimpeded. If a rootkit makes its way into your system, conventional software scanners may still go about their business in the normal manner — scanning memory, processes, and Registry hives, producing scan results that smugly claim, "no infection found." The operating system is changed or tricked by the rootkit into reporting false results. In the end, both the scanners and the users are deceived. We can help you see past a rootkit's trickery.

Rootkits not only hide themselves, they also hide their malware-associated processes, files, Registry entries (on Windows systems), and ports. This malware-hiding capability is what makes rootkits so dangerous — and it is their whole reason for being. A rootkit, in and of itself, does not present a danger — it just makes danger easier. It only becomes dangerous when it is used to conceal illicit activity — or if it is exploited by other malware programs that seek to conceal their presence.

No operating system is immune

Rootkits are very platform-specific. Although Windows systems are by far their most frequent targets, rootkits were first developed on Unix systems. That is where the term comes from: *root* (administrative) access and *kit* (a Unix break-in tool). Linux, of course, is a derivative of Unix — so it has its own (smaller) subset of rootkits. You should also know that Mac OS X has a rootkit on record (see www.theregister.co.uk/2004/10/25/mac_rootkit_opener/).

Typically, malware writers invest their time writing programs that attack whichever platforms

can reap them the most benefit — whether that means bragging rights (as in the early days) or illicit financial gain. No wonder so many malware programs are written for the popular Windows XP and Windows 2003 platforms — they get maximum exposure. Although malware writers usually won't waste their time writing for outdated Windows platforms or unpopular operating systems, any platform can attract their unwelcome attention by becoming more widely used.

But even though the rootkit serves to hide the activities and infected components installed on a system — as well as itself — all is not lost. Luckily, they have not yet reached the level of sophistication required to completely dupe all scanners. By understanding what rootkits are and how they work, you become better prepared to protect your computer or network from this security threat. The following sections explore these topics in more detail.

Why do rootkits exist?

As with many technological developments, rootkits have both good and bad uses. A rootkit by itself works like a hidden empty safe or vault. What matters is not the container itself, but whether it's ultimately used to store (so to speak) diamonds or vials of anthrax. A rootkit can hide a legitimate backup image of your operating system so your system can recover if it crashes — or the same little cache can tuck away a backdoor trojan. Although what's *in* a rootkit is of primary importance, there are ethical considerations at work. Legitimate uses for rootkits do exist — but many computer users oppose *any* use of a rootkit, regardless of whether its purpose is beneficial (whitehat) or malicious (blackhat). Some users object strenuously — and understandably — to anything being hidden from them on their own systems.

There is an even more compelling reason to object to including a rootkit of any kind — even a whitehat rootkit — in a program. Once a rootkit is known to exist, malware writers see it as an opportunity; They'll attempt to exploit its powers of concealment for their own benefit. Thus even whitehat rootkits pose a potential risk, which is why they're met with such criticism. A better technique is to employ encryption to ensure that critical data remains inaccessible and unaltered.



Any rootkit, regardless of its intended purpose, may be exploited by the bad guys to invisibly compromise a system.

All these efforts are aimed at hiding the presence of the intruder and the rootkit itself. Just as a thief who steals your wallet does not want to get caught, cyber criminals also try to maintain a low profile, so they can operate under a shroud of concealment.

Some deliver puppet masters

One common goal of a blackhat rootkit is to install a *puppet master* — to conceal a worm or trojan that takes over your computer and makes it a willing workhorse for malicious purposes. The usual technique is to hijack and

secretly maintain an open port that functions as a hidden backdoor, facilitating information transfer to and from your computer. Because the rootkit provides a shield of secrecy, such operations proceed stealthily and without interference. Your computer may have been recruited in such a manner to perform any (or all!) of the following dastardly deeds:



- ✔ **Launching Distributed Denial of Service attacks (DDoS) (or Night of the Cyber Dead):** The blackhat hacker may be recruiting your computer as a zombie or an unwitting accomplice to conduct a DDoS or Distributed Denial of Service attack on another system or network server.

The object of a DDoS attack is to bombard a system or network with so much traffic that it becomes inaccessible to legitimate users. Computers are normally recruited en masse to launch a successful DDoS attack — all without the users' knowledge. Broadband subscribers who have “always-on” connections are particularly vulnerable to becoming members of the cyber-zombie army. Successful DDoS attacks have been launched against Microsoft.com, Apple.com, Yahoo, eBay, Amazon.com, and the Million Dollar Homepage (www.milliondollarhomepage.com/), to name only a few.

- ✔ **Sending spam e-mail:** An infected computer may be used to launch e-mail spam attacks against targeted computers by sending out a multitude of solicitous e-mails. The zombie computer owner gets blamed for spamming, and the true source of the spam remains anonymous. Many zombie computer owners often have no idea their systems are being used for such illicit purposes — and their first wake-up call may come in the form of a letter from their Internet Service Provider (ISP) which threatens them with suspension of service for spamming.
- ✔ **Hosting and distributing illegal material:** A rootkit may be used to conceal the fact that your computer has been recruited to store and distribute illegal or pirated content. Such content might include music or video libraries, or even criminal pornographic materials. Storing the content on the hard drive of a recruited victim's computer kills two birds with one stone: It enables the true content provider to conserve on their own hardware resources, but more importantly it enables them to dispense criminal content with little risk of being identified or prosecuted. This is because the evidence resides on the compromised system not their own.

Some are just spies

Rootkits that act as spies enable *keyloggers* and *packet-sniffers* — programs that hide on a user's system and (respectively) log the user's every keystroke and inspect the data transmitted to or from the user's system or network — to do their dirty work. Privacy? Forget it. And it gets worse. . .

- ✔ **Breaking the bank:** Keystroke logs can be correlated with Web page visits to aid in the extraction of private and sensitive data such as bank login information, credit card numbers, and the like. This information can then be transferred remotely to the bad guys' computer and used to conduct criminal financial transactions or commit identity theft. A rootkit is an ideal hacking tool because it allows an intruder to maintain a connection that cannot be detected by the user. This enables data transfer to progress without interruption.
- ✔ **Harvesting your habits:** Another less insidious — but very annoying — form of spying is practiced by adware companies; at least one of them is known to employ a rootkit to prevent the removal of its software (if you can't find it, you can't remove it). The collection and transmission of information that reveals a user's browsing habits is very valuable to commercial adware companies. This type of spying allows the companies to serve up targeted pop-up advertisements that are custom-selected to appeal to the user. The now-discontinued Apropos rootkit (distributed by the adware company ContextPlus, Inc.) performs this function — and frequently churns out new variants to dodge current removal techniques. Just when we thought it was safe to go back in the water, a new and even more devious adware rootkit has emerged — as if to take the place of the retired Apropos rootkit. Certain variants of Link Optimizer adware can be installed by the Gromozon rootkit, which arrives via a WMF (Windows MetaFile) exploit (on unpatched computers). This infection is extremely difficult to remove, and utilizes other sneaky techniques besides rootkit technology, to ensure its survival. For more information on this threat, please refer to the following description provided by Symantec, entitled “Gromozon.com and Italian spaghetti”, and available at www.symantec.com/enterprise/security_response/weblog/2006/08/gromozoncom_and_italian_spaghe.html.
- ✔ **Sniffing the goods:** A *sniffer* is a common rootkit snooping tool that an intruder can install to capture all data transmitted over a network. Though network administrators may have legitimate uses for sniffers, a blackhat hacker uses a sniffer with a more devious intent. The captured data can be saved and analyzed to extract user login information. These stolen passwords are very valuable to an intruder, allowing an attacker to log on remotely and take anything the network has to offer — at the stolen password's privilege level. In this manner, an attacker can penetrate the network access files and retrieve all sorts of confidential and potentially valuable information.

Chapter 2

The Three Rs of Survivable Systems

In This Chapter

- ▶ Formulating resistance
 - ▶ Practicing recognition
 - ▶ Planning for recovery
-

A *survivable system* is a computer system that can survive a potentially constant onslaught of malware (and rootkits) on a regular basis. The three Rs in the title of this chapter describe that kind of system: *resistance* (being difficult for malware to attack); *recognition* (detecting and identifying malware when it's present); and *recovery* (bouncing back after a malware infection). This chapter is your first step toward gaining the knowledge and tools you need to protect yourself, your families, your networks, and your business from the real dangers of malware and rootkits.

But protection doesn't mean that you take only preventive measures. As computer hardware and software have become more sophisticated, so has the malware. As malware writers discover ways to sneak past existing security tools, malware can still find its way onto your system (although prevention greatly reduces your vulnerability). Recognizing this, we not only explain the basics of keeping malware out, but also offer tips that help you recognize problems when they exist — to recover from an infection and reset your computer to a happier state of being.

Note that this chapter is an overview of the three Rs. Parts II, III, and IV of this book describe resistance, recognition, and recovery (respectively) in greater detail.

Formulating Resistance

Just as lightning takes the path of least resistance and a chain is only as strong as its weakest link, cyber-criminals prey on systems that are the most vulnerable. Whether your computer becomes a sitting duck, inviting a malware attack, is mostly in your hands. If you don't want your computer or network compromised, then some precautions are necessary. That's truer than ever in today's sophisticated online environment.

In the past, protecting your system meant adding a decent firewall, as well as antivirus software — and not much else. But new malware programs are much more pernicious and deceptive than those from yesteryear. To protect your computer from becoming infected by rootkits and other malware, you need industrial-strength protection, more knowledge, and a greater understanding of your options.

Most of being secure is preventing attacks. Prevention means developing good habits and practices. As with defensive driving on the road, good computer security becomes second nature after a few months — and it can save a life (in this case, your computer's life). The preventive measures, programs, and procedures that we recommend here have been developed and tested by the experts. They really work. We have applied them on our own systems, and in the past five years Larry has had to deal with only six trojans and one virus — even though he is online every day doing intensive research and roaming all over the Internet. (Remember, without all the preventive steps described here, the number of problems would skyrocket.)

Hackers may not be smarter than you

Criminal hackers may be smart and experienced in computers, but it's unlikely they can stop themselves from making mistakes that can assist your investigation. Besides, if they were really and truly intelligent, they'd find a more lucrative endeavor than an e-B&E (*electronic break and enter*). North American jails are filled with B&E artists, many of whom were also smart but made mistakes, too. Criminal hackers as a group may not take kindly to these words, but the truth bytes!

You can deal easily with someone rummaging through your files by being prepared to respond to the threat in a manner which preserves your integrity and dignity. Much grief and heartache can be avoided if you are prepared ahead of time.

Learning to be a smart computer user is similar to taking swimming lessons: Learn the basic techniques, practice them often, be open to discovering more, and you will do well. You can start right away by making and applying a security plan that includes the elements we lay out in the following section.

Steps to a Better Security Posture

Three to five years ago, you could get by online with just a firewall and an antivirus program for protection. Much has changed since then. Now, even as the viruses continue to proliferate, we have many more forms of malware than ever — and they're even more insidious and destructive. Malware “surprises” are not only bundled with other programs you download, but they also come in via e-mail, IRC channels and other instant messaging services (such as AIM, Yahoo! Messenger, or Windows Messenger), corrupted Web sites, and zombie computers already infected by rootkits — to name only a few of the culprits.

Safeguarding your computer against rootkits is essentially the same as preventing malware such as trojans, worms, and backdoors from gaining a foothold. Rootkits need to piggy-back with these other types of malware in order to get installed on your computer in the first place. No protection is 100 percent effective, of course, but you can — with a little help — get closer to that state of bliss. As more people begin safeguarding their computers more effectively, successful malware attacks will be reduced as the bad stuff has fewer opportunities to spread. But by the same token, malware propagation is also a factor of clever malware writers continually trying to outwit those who are in the business of removing infections. It is unlikely that this cat-and-mouse game will ever cease. As soon as the good guys develop a solution to “L' Infection du Jour” (The Infection of the Day), the bad guys will be at it trying to circumvent their approach. It is a challenge of wits — and some say the malware writers always have the upper hand because they don't have to worry about ruining their reputation or wrecking your system if their creations backfire.

The following list gives you the basic synopsis of what the security gurus have been trying to help you with in relation to your online activities for the past several years. The following list (and the added detail on these topics later in the book) can go a long way to helping you prevent malware and rootkits from infecting your system.

- ✔ **Use limited-access user accounts.** Provide user accounts with limited privileges, and use them instead of the Administrator account when you go online. Malware has access to your computer according to the privileges in the account you are using — the more privileges the account has, the more risk you run. See Chapter 4 for more information on limited-access user accounts.

- ✔ **Set local or group security policies.** Security policies establish what incoming stuff is accepted or rejected when you're online — or what specific rules and practices cover the users on your computer or network. Full instructions are provided on setting up your security policies in Chapter 6.
- ✔ **Configure and enable security logs and auditing policies.** Auditing on a Windows XP Professional computer helps you monitor what happens on it, providing greater security and reliable information for troubleshooting problems. In Windows XP Home Edition, the Event Logging service is turned on automatically by default, but you can't configure it. Event Logging will record all Application, Security, and System Events and no selective auditing procedure is required. In Windows XP Professional, auditing is turned off by default — but you can configure it exactly as you want it. Chapter 6 explains auditing in more detail and Chapter 8 will fill you in on some basics of inspecting and interpreting event logs.
- ✔ **Use a firewall.** A *firewall* is a program that puts up a protective logical barrier between your network (or individual computer) and the outside world (the Internet). It monitors Internet traffic to thwart unsolicited intrusions — its function is to let the good stuff in and keep the bad guys out. A firewall is essential to shielding both servers and individual computers. On networks, use both server and workstation firewalls. (See Chapter 4 for more information about firewalls.)
- ✔ **Use scanners.** Use scanners on all systems — including antivirus, anti-trojan, anti-spyware, and anti-malware applications. *Scanners* are programs that look for evidence of infection on your hard drive (or other specified storage devices), and disinfect or quarantine it when found. Scanners typically look for infected files, folders, and Registry entries. Some scanners can also scan memory for evidence of infection. (See Chapter 4 for more information on the various types of scanners and the features you need to look for.) Keep them updated and make sure their active protection (if it applies) is enabled on startup.
- ✔ **Use an alternate Internet browser.** Internet Explorer (IE) Version 6, although much improved over previous versions, is very old. The original Internet Explorer was introduced before anyone had any clue that the hackers-in-waiting were ready to explore and exploit every one of its nooks and crannies. Malware creators have had plenty of time to pick it apart and find its weaknesses. Though its vulnerabilities have been regularly patched, and it has undergone significant improvements, it still requires configuring to achieve better security. In Chapter 4 we show you how to harden your security with Internet Explorer. You may also want to investigate more secure browser solutions such as Opera, Mozilla Firefox, or IE7 for regular surfing online. We briefly discuss the enhancements introduced by Firefox, Opera, and IE7 in Chapter 4.

- ✔ **Securely configure your Web browser.** Internet Explorer versions 6 and prior are not secure in default form. A lot of malware can be stopped by following the instructions for configuring IE securely (and you can find those in Chapter 4).
- ✔ **Create strong, unique passwords and store them with encryption.** Never use the same passwords over and over. Change them regularly. Using strong account passwords will prevent hackers from gaining access — both online and physical — to your data. Using encryption applications (provided on this book's CD) will help you to store and remember your passwords safely. Flip to Chapter 4 to learn more.
- ✔ **Stay current with automatic Windows updates and security patches.** Windows and Microsoft updates and security patches are now more important than ever. Malware creators know that most people are reluctant to put up with the hassle of getting updates installed, so they gobble up information about vulnerabilities and (to paraphrase Maxwell Smart) use it for nastiness. Setting your operating system to automatically download and install Windows Updates will relieve you of the burden of remembering to perform this all-important task. For detailed coverage of patches and updates, see Chapter 5.
- ✔ **Scan regularly with anti-rootkit tools.** Technically speaking, all of your anti-malware tools — including your antivirus and anti-spyware applications — are indirectly anti-rootkit tools, too. They help prevent the means by which rootkits get introduced to your computer. There are also, however, tools to use specifically against rootkits — they're explained in Chapter 9.
- ✔ **Apply an updated HOSTS file.** The HOSTS file has no extension and resides in your WINDOWS folder. It can help you filter the addresses of bad Web sites so your browser will be prevented from navigating to them (the http address will resolve to your own computer). See Chapter 4 for more information on the HOSTS file.
- ✔ **Maintain regular full-system and file backups on removable media.** Ensure the use of — and access to — backup software are protected with strong, rotating (that is, frequently replaced) passwords.
- ✔ **Encrypt personal account files.** Another feature provided by Windows XP Professional is the capability to encrypt your account files so no one else can read them. You can learn more about EFS by entering those three letters into the search box in your Help and Support section in the Start menu. You can also read more about it at the following URL:

www.practicalpc.co.uk/computing/windows/xpencrypt1.htm
- ✔ **On networks, implement both Intrusion Detection (IDS) and Intrusion Prevention (IPS) Systems.** Intrusion Detection Systems (IDS) protect networks and host computers by monitoring the data communications

between them to discover possible threats. An IDS is like a guard who stands outside the door, checking anyone's credentials before allowing them access. Intrusion Prevention Systems, on the other hand, act as real-time monitors and resident scanners, ready to remove any intruders. IPS is like the guard on the inside of the door, kicking out anyone who is not allowed, should they get past the first guard. Please refer to Chapter 6 for more detailed information.

- ✔ **Harden your Windows operating system to improve security.** Windows XP comes with a large number of services that run in the background. Many of these services are nonessential or can be set to start manually on an as-needed basis. Some have been implicated with exploitable vulnerabilities, though patches have been released to remove those threats. We show you how and what to disable in Chapter 3.



Some of the best security tools, applications, and methods are available to you on the DART (Dummies Anti-Rootkit Toolkit) CD, to help protect you and your computer from this scourge of malware. This book's Appendix outlines all the tools the CD has to offer.



Protecting your system from such threats is not solely a function of installing the latest state-of-the-art security programs. Practicing common-sense surfing (which the ultra-hip among computer users call “safe hex”) is critical to ensuring the continued safety of your computer or network. Network administrators must take proactive measures to fortify the host, protect all connected client computers, and to educate their users about developing good security habits.

We will show you techniques, methods, and applications based on these essential criteria in later chapters. Read on.

Practicing Recognition

If you have ever been infected by malware, you most likely remember what your computer was doing — before, during, and after — but if other users (particularly kids) have used your computer, in all likelihood they will not have a clue what happened — or at least won't provide you with one. Because the symptoms of malware can be so diverse, even if you have had a “malware experience,” you probably are not familiar with all of them. In this section we will acquaint you with the most common malware symptoms, so you can identify their potential presence. But before you read further and embark on diagnosing your system, keep the following points in mind:

If you're hacked . . .

We can advise you not to take it personally if you're hacked — after all, to the typical hacker, you're probably a total stranger — but how you respond depends on you. If you are prepared to deal with this form of emergency, you'll make a

lot fewer mistakes. People who are unprepared often “panic” or fly into fits of rage, neither of which solves the problem. Computers are logical; irrational feelings will only lead to more mistakes and greater difficulties.



- ✔ Compiling a complete list of malware symptoms — today or in the future — would be impossible. The symptoms we describe in the next section are presently the most common as we write this book. Use this list as a guide to recognizing the changes that malware can effect on your computer.
- ✔ If you have a rootkit, your system may not exhibit any symptoms at all.
- ✔ Interpreting the cause of symptoms can be tricky. One symptom may indicate you have a malware problem — or it could be a simple hardware problem. Other malware-type symptoms may develop when programs do not play nicely together, or too many programs run simultaneously (many such programs are nonessential but are set to run automatically at startup). In that case, the problem may not be malware at all, as discussed in the “Recognizing when the problem isn't malware” section later in this chapter.

Spotting signs of malware

The operational symptoms presented if a rootkit is hidden on your system will not substantially differ from those experienced if a “normal” (non-rootkit) infection is present. That is because the rootkit's entire purpose is to hide the presence of traditional types of malware. Your computer will therefore react the same way as it would if it were to become infected by the same threat without rootkit concealment. However, since a rootkit is often used to establish remote control over the target system, insidious backdoor threats are more frequently encountered when rootkits are involved. The following system symptoms often indicate the presence of malware:

- ✔ Prolonged startup or shutdown times.
- ✔ Sluggish performance, system crashes or hangs, and blue screens of death (BSOD).

- ✔ Internet-connection attempts or firewall alerts that indicate unknown processes or programs are requesting Internet access or “phoning home.”
- ✔ “Hijacking” of browser startups and of search pages that take you to suspicious sites you have not chosen to visit.
- ✔ Newly installed — but suspicious-looking — toolbars, browser favorites, or desktop or taskbar shortcuts appear out of nowhere.
- ✔ Bogus Windows security alerts redirect you to a suspect “security” vendor’s Web site in an attempt to sell you a program.
- ✔ New, unknown startups appear in MSConfig (the Windows system-configuration utility), or new unknown processes appear in Task Manager.
- ✔ You cannot open the MSConfig, Task Manager, or Regedit programs.
- ✔ New, suspicious, non-Microsoft services run in the Services console.
- ✔ A hard drive is often too “busy” (churning away) for no apparent reason.
- ✔ Task Manager indicates that CPU consumption is near its maximum.
- ✔ Inexplicable port activities.
- ✔ Desktop wallpapers change to infection warnings.
- ✔ Programs unexpectedly close when you try to open them (especially Regedit, Task Manager, and MSConfig).
- ✔ Your home page or search pages have been changed in your Web browser and you cannot change them back.
- ✔ Web-browser Favorites or Bookmarks have mysterious new additions that you didn’t put there.
- ✔ In Internet Explorer, Internet Options has vanished from the Tools Menu and from the Control Panel on your computer.
- ✔ You are blocked from surfing to anti-malware and security Web sites but not from accessing other sites.
- ✔ You are plagued by persistent pop-up and pornographic ads, even offline — and often you can’t close them.
- ✔ Anti-malware programs — or your firewall — will not open or load at startup.
- ✔ You find new icons on your desktop or in System tray.
- ✔ Strange message windows appear on the desktop, displaying phrases such as *Your PC has been Nuked!* or *I am in control of your computer now, hahahahahahaha!*

Recognizing when the problem isn't malware

Malware has (understandably) received a lot of publicity, so people tend to assume that if they have problems with their software or computers, the problems must be due to a malware invasion. Actually, we've found — while helping folks who claim to have malware problems — that the trouble is often due to annoying but relatively “normal” problems in the operating system or hardware.



Not all problems that you may encounter with your Web browser and operating system are due to the presence of malicious software. There's a lot of it out there, but other practical explanations may be more on target.

If you continue to have problems after running numerous anti-malware scans — that haven't found anything wrong — go to one of the security Web sites listed in Chapter 13 and ask for assistance. Alternatively, you can take your computer to a reputable repair shop for diagnosis and repair.

Suspecting that you've been compromised

If you have a rootkit, you may or may not notice any symptoms. That depends on its payload of malware. Chapter 8 details how to recognize the tracks that rootkits may create and leave behind.

Planning for Recovery

Do you know how to ensure that your computer data, privacy, and sanity stay intact if the worst-case scenario — infection by a rootkit — occurs? Well, here's a quick checklist:

- ✓ Before an attack happens, establish that your computer is clean and secure. Scan it with everything you can, double-check your security policies and applications, ensure that all performance maintenance tasks are done, and make sure it's fully updated. You may wish to get help from the experts for these tasks. Please refer to the chapters in Part II of this book for more information on these chores.
- ✓ When you know you have a clean machine, make a backup of all of your files — and keep those files up to date.

- ✔ Be sure you have the Install disks for your operating system — on the removable medium appropriate to your system (such as CD-ROM, floppy disk, or removable drive).
- ✔ Create a back-up image of your hard drive using a program like Acronis Imaging, which is included on your DART CD-ROM. Acronis Imaging can quickly restore the content of your hard drive in one fell swoop with its Snap Restore feature. It is best if this backup image is stored on another medium such as an external hard drive to ensure its cleanliness. See Chapter 4 for more information on backing up your files, and see the Appendix and Bonus Chapter 2 for more information on Acronis Imaging.
- ✔ Set a new Restore point after each full backup is completed. On Windows XP, having a Restore point can be a Godsend if you run into trouble. You can use the point to restore your computer to a clean slate and use the backup to (hopefully) replace any files you may have lost. See Chapter 4 for more information about restore points.
- ✔ Have hard copies available of complete instructions on how to wipe your hard drive, reformat the drive, and reinstall your operating system. See Chapter 11 for more information.
- ✔ Keep backup copies of your cherished (work-critical or personally important) programs and security applications. Date and store them in a safe and secure place, preferably a locked cabinet in another room. Make frequent backups of your Documents and Settings folder, as this folder changes daily on most computers and you stand to lose a lot if you don't.


These are some of the basic tasks that can make you an effective rootkit survivor. We guide you through them, step by step, in Chapters 10 and 11.

Part II

Resistance Is NOT Futile

The 5th Wave

By Rich Tennant



In this part . . .

We show you how to beat the enemy forces at their own game. Before heading out onto the Internet battlefields, you need to be prepared, fully equipped, and ready. Here's where you train for the conflict — hardening your computer, toughening network security, sharpening your vigilance, using an array of methods and equipment to protect yourself wisely.

Malware enemies are like the Borg in the *Star Trek* universe — they want to add your distinctive attributes and technology to their own. You don't have to let them. Resistance is *fertile* — bringing victory and “QuPLA!” (That's Klingon for “Success!”)

Chapter 3

Practicing Good Computer Hygiene

In This Chapter

- ▶ Setting a restore point in case of emergency
 - ▶ Cleaning up files and programs for a more secure system
 - ▶ Revising your Startup List
 - ▶ Learning to use Windows system-maintenance tools
-

Housekeeping is a necessary chore that helps maintain your health and well-being. As with your house, life is just better if you maintain a tidy and well-run computer. Some of the most frequent computer problems are system slowdowns, corruptions, and crashes stemming from useless junk files — and malware — gathered over the course of months (or even years). An unkempt computer increases the difficulties encountered when attempting to diagnose and remedy the presence of rootkits — they're even harder to find if they have a big pile of junk to hide in. We can help — so fire up your virtual vacuum cleaners, and let's clean up!

Most of your computer maintenance is done within Windows itself, but for the rest, a few applications will be needed. Fortunately, the applications featured here are reputable, tried-and-true tools that have been tested by the experts.

Before Doing Anything. . .

Before touching a hair on Windows' pretty head, you should do a few things first: Get friendly with System Restore, back up your Registry, run the Windows Backup utility, and pare down your automatic startups with an eye toward improving security. The idea here is to make sure you can get your computer up and running again, no worse than it was and with no data loss, if a change you make goes horribly wrong. But first things first: Set a restore point. This section shows how.

Using System Restore

Who hasn't made mistakes installing new software or hardware that suddenly makes the computer go haywire? Similar to the Undo command in many applications, System Restore allows you to turn back the clock, undoing the actions that caused your system to have problems.

It can save your skin.

The following sections describe System Restore and show you how to access the utility and how to use it to turn your machine back to a time when it worked properly if you need to.



You need Administrator privileges to use System Restore, but automatic restore points are made without regard to who is using the computer.

Knowing how System Restore works

System Restore preserves the state of your computer by archiving copies of critical system components such as your Registry, local profiles, and important system files and settings on a regular basis. Each of these *restore points* allows you to time-travel (in a sense) from a non-functional system back to a time when it worked properly.

You can create your own restore points whenever you want; the time and date for each one will automatically be recorded. Depending on how often you use your computer, you may have a few or many restore points. By default, Windows creates a restore point if 24 hours have passed since your computer was last used — and also in response to significant events such as installing an unsigned device driver, installing Microsoft's Automatic Updates, performing a System Restore operation, restoring data from a backup, and so on.



Windows maintains restore points for up to 90 days, after which they're discarded. System Restore will not alter your personal data, such as Word documents, graphics files, browser favorites, or game saves, and it won't uninstall applications, either.



One of the authors needed to use it recently after attempting to install new hardware (a wireless Ethernet card) on his personal business laptop. He loaded the drivers and the software to run it, added the hardware and restarted. The Welcome screen didn't appear, and Windows XP would not accept his administrator logon. Safe mode would not work. Oops! With assistance from the manufacturer's tech support, the hardware was removed. After multiple reboot and logon attempts, it did finally accept the logon. He immediately uninstalled the drivers and software, then applied a restore point from at least a week before. Phew! That was a close one. The system was restored to a stable condition again. Thank you, Microsoft, for this lifesaver!

Accessing System Restore

You can access the System Restore Wizard in one of three ways: from the All Programs menu, via the Help and Support option in the Windows XP Start menu, or from the command line.



Before using System Restore, you should save any changes and close any applications you have been using — because it requires you to do a system restart.

- ✓ From the All Programs folder, click Start and choose All Programs ⇨ Accessories ⇨ System Tools ⇨ System Restore.
- ✓ Via Help and Support, Click Start and choose Help and Support. Under Pick a Task on the right side of the window, find and click *Undo changes to your computer with System Restore*.
- ✓ From the command line, click Start, choose Run, type **cmd** into the Open: field, and press Enter. When the command prompt appears, enter the following command to launch system restore: **C:\WINDOWS\system32\restore\rstrui.exe**. (Where C: refers to your primary drive where Windows is installed.)

Creating a fresh restore point

You may want to create restore points yourself before using Regedit to edit the registry, or if you're applying a new configuration to your computer by changing the policies or services settings (see Chapter 6). You can name them anything you want, and make them at any time. To do so, follow these steps:

- 1. Start the System Restore Wizard as described in the last section.**
- 2. Select the Create a Restore Point option on the right and click Next.**
- 3. Type in a name for your restore point and click Create.**

System Restore will automatically assign the date and time when you click Create.

Your new restore point is a reality, and will remain in System Restore for ninety days before it's deleted.

Changing the Drive Settings

You can make changes to how much drive space System Restore uses by clicking System Restore Settings on the left side of the Welcome screen. The System Properties dialog box appears.



You may also access the System Restore settings by right-clicking My Computer, selecting Properties, clicking the System Restore tab, selecting the drive you want on the left, and then clicking the Settings button.

On systems with hard drives larger than 4 GB, System Restore uses up to 12 percent of the space to store restore points. On systems with less than 4 GB, it uses 400MB.



If you're trying to increase hard drive space, you can reduce the amount of space System Restore uses, but be aware that this will result in fewer restore points.

To change the amount of space System Restore uses, move the slider to the left. The less drive space allowed, the fewer will be the restore points available. Move the slider to the right to increase it if the slider isn't all the way to the right.

Note that you can also turn off System Restore for the selected drive by clicking the Turn Off System Restore on All Drives check box on the System Restore tab.



This procedure erases all your restore points on that drive.

Click OK when you're done configuring System Restore.

Clearing out System Restore or shutting it off



When cleaning malware and rootkits, there are instances when you may need to turn off System Restore, since the malware may have corrupted the system restore data. Normally, after your system is completely clean, if malware files remain in the system volume information, then all System Restore points are flushed. *Flushing* just means you turn System Restore off, restart, and then turn System Restore back on again. This will effectively eliminate (flush out) any malware that may reside in the System Restore files. Flushing your System Restore points helps prevent infected material from being restored to your system in the event that you must perform a system restore at a future date. Following this procedure, it is necessary to create a new restore point. Though this procedure is not very complicated, we suggest that you only do this with the advice and guidance of a computer expert. This technique is used primarily for guided system cleanings. It not only flushes the Restore points, but all the data contained in the system information volume which System Restore uses. If something goes wrong, and you are unable to use System Restore afterwards, you will need expert help.



System Restore protects critical system files from changes or deletions by copying them before the changes take place. As can sometimes happen with antivirus or anti-spyware software, before it can quarantine or delete an infected file, System Restore will copy it for “safe” keeping. After your scanner cleans your computer, the file may be clean on the rest of the computer, yet show up on an antivirus or anti-spyware scan as being in the System Volume Information. At this point you may want to clear out the System Restore Points. Simply turn it off, close the box, then open it again and turn it back on. Then set a fresh Restore point.

Restoring from a restore point

If your system fails completely, press F8 during Startup and choose to restore from the last good configuration (probably the one in effect when you last set your restore point). If you can boot into Windows but need to restore anyway, run System Restore and choose a point to restore from.

- 1. Start System Restore using one of the methods described in the “Accessing System Restore” section earlier in this chapter.**

The Welcome to System Restore window appears.

- 2. Click the Restore My Computer to an Earlier Time option on the right.**

The Select a Restore Point screen appears.

- 3. Choose a date in bold on the left and then click to choose a restore point on the right.**

- 4. Confirm the restore point you have chosen and click Next.**

System Restore restores your computer to the restore point and then restarts your computer.

Running System Restore from Safe mode in an emergency

It is possible that you may be unable to run system restore in normal mode. If that happens, and your computer is still bootable, you should try to run it in Safe mode from the command prompt, by following these steps:

- 1. Restart or boot your computer, tapping the F8 key while it loads.**

The Windows Advanced Options Menu appears.

Some models of computers may give you an error message when you do this. If so, restart the system and try again.

- 2. Select the Safe mode with Command Prompt option.**

3. **Log on using your normal user profile name if it has administrative privileges (or log on as the Administrator if it does not).**
4. **At the command prompt, enter the following command to launch System Restore:**

```
C:\WINDOWS\system32\restore\rstrui.exe
```

Here C: refers to your primary drive on which Windows is installed.

5. **Press Enter.**

The System Restore wizard appears. Its prompts will guide you through System Restore procedure.

Backing up your Registry

The *Registry* is your Windows computer's marching orders. Before it does anything, the computer checks the Registry for instructions on what to do and when to proceed. Changes occur to the Registry whenever you use your computer, so before you proceed with any cleaning or file-removal tasks, you should back it up. Should anything go horribly wrong, you can restore the original settings and Registry so you can try again.

There are four primary methods you can use to back up your entire Registry and system settings on Windows XP:

- ✓ The Registry Editor's Export function (Regedit)
- ✓ System Restore (see the "Using System Restore" section earlier in the chapter)
- ✓ The Windows XP Backup Utility (see the next section)
- ✓ Third-party backup software.

Included on the DART CD are two different third-party backup utilities: The Replicator and Acronis True Image. The Replicator is versatile, allowing you to make backups to almost any kind of media, and even across networks. Acronis True Image makes high-quality image backups to DVD-RW media or to an external hard drive. Backup tasks for both of these programs can be automated.



Using the Registry Editor to back up the whole Registry

Even for those of you who feel comfortable editing the Registry, backing up the Registry is a *must* before making any changes. By making a backup copy of your Registry, you can recover from any changes made in error.

Though using Regedit is the easiest method, Microsoft recommends using the Windows XP Backup Utility or System Restore as the preferred method, to back up and restore the entire Registry.

To back up your Registry by itself, follow these steps:

- 1. Open the Registry Editor by clicking start and typing regedit into the Open: field and click OK.**

The Registry Editor appears.

- 2. Choose File ⇨ Export.**

- 3. Specify a name for your Registry backup file, choose a location for the backup file, and click Save.**

Registry Editor automatically saves as a `.reg` file using the filename and location you have provided.

Using the Registry Editor to back up a key or a branch of the Registry

If you're editing the Registry and making only minor changes, you should only back up the key or section that you intend to make changes to prior to editing it. By restricting the backup data to the section you're changing, any introduced errors can be more easily corrected. All subkeys in the indented sections below the saved key will also be backed up.

In the following example, we illustrate how to back up the WinLogon Registry key:

- 1. Open the Registry Editor by clicking start and typing regedit into the Open: field and click OK.**

The Registry Editor appears.

- 2. Navigate to the key you want to change, in this case we have selected WinLogon.**

- 3. Right-click the key and select Export from the context menu.**

- 4. Type in the name of the key you're saving (in this case, Winlogon) and click Save.**

The `.REG` extension will automatically be appended to the file.

Restoring a previous Registry or key

If you need to restore the Registry or a Registry key from a backup file, browse to the backup `.reg` file, double-click it, and then click Yes when asked if you want to add the information to the Registry.

Backing up your stuff with Windows Backup

Before you start deleting those old programs and files, you should make backups of anything you want to keep. For this, you may want to use the Windows Backup utility (included with all Windows versions) or a commercial backup product.

The Windows XP Backup Utility is not compatible with CD/DVD-RW drives. You will need to use almost any other kind of media — such as a Zip drive, USB or FireWire memory stick, external hard drive, or a separate partition on your internal hard drive. You can find more info in a Microsoft online document at

www.microsoft.com/windowsxp/using/setup/learnmore/bott_03july14.msp

Installing the Windows XP Backup Utility

Note that by default, the Backup Utility is not installed on Windows XP Home Edition, although it is preinstalled on Windows XP Professional Edition. For Windows XP Home Edition, you need to install the Backup Utility manually, like this:

- 1. Insert the Windows XP disc into your CD drive.**
- 2. Browse to the following location:**
`D:\VALUEADD\MSFT\NTBACKUP`
where D: is the letter of your CD drive.
- 3. Double-click the NTBACKUP.MSI file to start the installation.**

Backing up your files with the Windows XP Backup Utility

The easiest way to use the Windows XP Backup Utility is to use the wizard rather than Advanced mode. Unless you decide to choose what to back up, this tool will also back up your Registry and system data settings. To do so, follow these steps:

- 1. Start the utility by clicking Start and choosing All Programs ⇨ Accessories ⇨ System Tools ⇨ Backup.**
The Backup or Restore Wizard appears.
- 2. Make sure a check mark is in the box beside Always Start in Wizard Mode and click Next.**

You will be prompted with the question `Will this be for a backup, or a restore?`

3. Select Backup Files and Settings and click Next.

A screen appears where you can choose what you want to back up. If you choose any of the first three items, the backup will also include your Registry and system settings. It will not do so automatically if you specify that you want to `Let me choose what to back up`.

You'll notice that throughout this wizard you always have the option to go back and correct or redo items, or to cancel at any time.

4. Select My Documents and Settings for this example and click Next.

5. Click Browse to select a drive on which to save your backup, type in the name of your backup, and click Next.

A summary of the choices you made for the backup so far appears. You can proceed with the backup operation from this point if you want, but there's more yet, read on.

6. If you click the Advanced button in the Completing the Backup or Restore Wizard window above, you can make several choices about what kind of backup you are making; Normal is good for a first backup, so choose Normal and click Next.

Here's what each backup option means:

- **Normal:** Copies all selected files and folders, and marks them as backed up.
- **Copy:** By selecting this option the files and folders you have chosen will simply be copied, but will not be marked as having been backed up. By using this with files you select yourself, you can make copies of many different files and folders, perhaps to share with others.
- **Incremental:** Only copies the files made or changed since the last backup. It marks them as having been backed up. If you use this feature you will need both the incremental and the normal backups in order to restore the files. It saves time and disk space; no need to re-copy everything all over again. You can use this to keep your normal backups up to date.
- **Differential:** It copies only the files made or changed since the last normal or incremental backup, but does not mark them as backed up. If you want to restore the files you will need the last normal backup as well.
- **Daily:** Copies the files that have changed that day only. It does not mark them as backed up. You might use this as part of a daily report that you make of particular files.



7. **You can click Finish to begin the backup, or you can continue with the Advanced options. You can also choose to verify, compress, or prevent shadow copies with your backup.**



We don't advise that you disable the shadow copy feature. With it on, it enables you to use your computer for other tasks while the backup program works in the background. You can even be using a file that's being backed up. Otherwise, if you turn it off, you will not be able to do anything else on your computer until the backup operation is complete. You might want to choose to disable it if you have malware operating on your computer, and the files being backed up are particularly important.

8. **In the next window, click Normal and then click Next.**
9. **Now make the choices according to the information provided in Step 7; For this example choose Verify Data After Backup and click Next.**
10. **In the next window, choose whether to overwrite or restrict access to your backup and click Next.**
11. **In the following window, choose whether to back up now or later and then click Next.**

A window appears summarizing your choices so far.

12. **To begin the backup, click Finish.**

Cleaning Your Windows to Improve Security

Hard drives get cluttered up with useless junk files, old firewall log entries, software installers, and temporary Internet files in Internet Explorer. Computers can be like a closet that accumulates loads of stuff — a lot of it useful. Sometimes you take stuff out and then later you put stuff back in. A lot of things might end up at the back of the closet — and you don't use them anymore but you don't want to throw 'em away, or it may be too much trouble to rummage around in there. And then a time comes when you open your closet door and all this junk falls out on top of you — leaving you no choice but to clean the closet. Computers are likewise prone to junk buildup.



One of the problems with letting your computer get all cluttered up with old programs and junk files is a too-full hard drive. The more memory and processor time used to find files on the hard drive, the less your system has to use for more important tasks. This can result in a real system slow down that you may even mistake as a symptom of infection. On most Windows systems, at least a third of the hard drive should remain open to act as a buffer for the programs that you run. If the hard drive gets too full, your computer will become sluggish and slow.

Everything and the kitchen sink: Loading only what you need at startup

Do you remember those fellow students back in school who used to raise their hands and actually wave 'em around when the teacher would ask a question, yelling, "Pick me, pick me!"? Well, unfortunately, just about every software application is just like those overeager students: They want your attention immediately when you start Windows. When installing new software, applications have a tendency to place themselves in the Startup folder or add themselves to the Registry autostarts so they'll load right into memory when Windows starts — and they even add an icon to the system tray in the bottom-right corner of your screen. Supposedly the applications do this so you can access them right away if you need them. But why is it so important for *every* new program to load as soon as you boot? When you get home from work at the end of the day, do you turn on all your lights and every appliance in the house, whether you need it or not? Of course not!

Seriously, how often do you really *need* all these programs? Take Apple QuickTime, for example, or Adobe's Reader: do you ever seek out and start these programs? Or do you only start them when you come across a file that requires them? With a little careful review, you can probably think of at least one application you don't use very often that still insists on starting when Windows starts.

The point is that no matter how convenient it may be to have all of these applications instantly available, running too many applications at startup can seriously impair your system's performance and stability. Each application requires a specific portion of your RAM and processor time to load and remain active. The resources provided to these unnecessary applications, prevent your system from using them for other tasks. The result can be a slow system that takes a long time to boot, shut down, and perform simple processing tasks. You might even suspect malware is the culprit, and run countless anti-malware scans, only to find that the problems persist.



Computers, for all their complexities, are actually glorified calculators. They do one operation at a time, even if at lightning speed. Indeed, although you may have numerous applications running in the background, your computer concentrates on only one of them at a time, one after another, depending on your requests and how you've configured them to work. Applications can interfere with each other, squabbling like children attempting to get a parent's (your computer's) attention — all at the same time and each for different reasons. Each new request calls on the remaining system resources to do more. If just enough resources are available to complete each task, the parent (in this case, your system) will get them all done — but very slowly. If the parent can't respond right away because of another child's demands, the other kids may throw a tantrum (crash) — or the parent may have a nervous breakdown (the system may lock up or crash).

To mitigate this situation, you can review the applications you have that load at Startup. Besides the BIOS and your operating system, you need only some security applications and your browser to load at bootup. The rest can be opened as you need them from shortcuts you can create. If you plan to be online, it's absolutely necessary to load a firewall, and the active protection components of your antivirus and anti-spyware applications. Other than that, what you load at startup is up to you — not the application vendors.

Many applications will have an entry in their configuration usually listed under Options or Setup that says something like `Load at Startup`. It's far easier and safer to disable this feature within the programs themselves, rather than by removing them from the startup list. Some programs, such as Apple Quicktime, iTunes, Skype, and RealPlayer, either do not have this feature or choose to ignore it, so programs of this type must be unchecked in the startup list.



Make a backup of your Registry before you make *any* configuration changes. Setting a Fresh Restore Point (in the section of that name earlier in the chapter) will backup your Registry and your important system settings. See the “Backing up your Registry” section earlier in the chapter.

Editing your startup list using MSCONFIG

To review your Startup applications using MSCONFIG, follow these steps:

1. Click the Start button and then click Run.

The Run box appears.

2. Type `msconfig` in the Run box and click OK.

The System Configuration utility appears.

3. Click the Startup tab.

A list of applications with check boxes beside them appears. The checked ones are set to load at startup.

4. Check your list against a Windows Startup Database, to see what you really need to keep.

Visit this URL to view the CastleCops searchable Startup List:

```
http://castlecops.com/StartupList.html
```

Alternatively, consult the The Bleeping Computer Startup Database (another useful resource of the same type) at

```
www.bleepingcomputer.com/startups/
```

By following the advice provided in the Startup Lists, you should now have some idea of what applications you need to keep.



5. In MSConfig, uncheck the boxes beside any entries that you have determined you don't need, and then click OK.



When you uncheck boxes next to Startup entries, be sure to do them one at a time — uncheck one box, reboot your computer, and then uncheck the next box and reboot again, until all the programs you don't want running at startup are unchecked. Yes, it's a hassle, but there's a good reason to put the time in: If you make one mistake, you can easily correct it — but if you try to uncheck them all at once and you make a mistake, you may get a huge mess that's difficult to correct.

Editing your startup list using Autoruns

As an alternative to the Windows System Configuration utility (MSConfig), you may want to use one of the many available third-party startup managers. One of the best of the bunch is Autoruns by Sysinternals.



Autoruns is included on your DART CD. See the Appendix for information on how to install it, and see Chapters 9 and 14 for information on using it.

Autoruns is superior to using MSConfig to manage your startups because it monitors many more locations than MSConfig does, including your startup folders. Because some of these locations are not well known (and are difficult to see), they have become preferred targets for malware.

When Autoruns is first started, the Everything tab is selected by default, and this will show you just what it says — all of your startups. However, you may click one of the 12 other tabs to concentrate on startup of a specific type, such as Drivers or BHOs (Browser Helper Objects).



There are a few things you can do to safely pare down the number of entries you have to wade through:

- ✓ **Choose Options ⇨ Hide Signed Microsoft Entries, and then hit F5 (to refresh the screen).** This safely eliminates Microsoft startups from the listing.
- ✓ **Click Options and uncheck Include Empty Sections.** Doing so lists only those locations for which startups exist on your system; empty sections are excluded.
- ✓ **Right-click an entry and select Google.** This automatically opens your browser to Google so you can investigate the startup program name listed under the Image Path column.

Each startup entry has a box next to it; unchecking that box deactivates the startup but leaves it in the Registry. Should you deactivate an item and change your mind later, you can simply recheck the item to reactivate the

program. If the program is malware-related, then you would want to axe it altogether by right-clicking the item and choosing Delete. This removes the autostart from the Registry so you don't have to edit the Registry manually.

Autoruns has many more powerful features; we show you how to use it to detect and remove rootkit startups in Chapter 9.

Removing unused programs

We can start by removing any user-installed applications you don't currently use — or don't expect to use at all in the future. First, check to see whether the program has an entry in the Add/Remove Programs page of the Control Panel. If it does, uninstall it from there — and then delete any unnecessary remaining folders or files afterward using Windows Explorer.

If the program is not listed in Add/Remove Programs and does not have its own uninstaller, then most of the time it's okay simply to delete the folder. To delete a folder with Windows Explorer, simply highlight it in the left pane, and then choose File ⇨ Delete. A dialog box pops up, asking whether you want this to go into the Recycle Bin. Click Yes and it's outta here.



Make a backup copy of your Registry or set a Restore Point (if you're running Windows XP) before removing any programs. (See the “Before Doing Anything” section earlier in the chapter for the skinny on how to do so.) Uninstall application programs one at a time, rebooting after each removal. (Yep, it takes time. But it's a whole lot safer.)

- 1. To access the Add or Remove Programs feature of the Control Panel on Windows XP, click Start and choose Control Panel, and then click Add or Remove Programs.**

A list of programs appears.

- 2. Examine each program in the list.**

Windows XP shows you how often you've used the program and when you last accessed it. This information may not be provided on other versions of Windows.



If you don't know what a particular program is — or does — *do not* remove it until you do. Lots of Windows system files, folders, and must-have-or-it-won't-work programs have strange names. If you can't find a particular program in the Startup Index — or at the links we list here — browse to CastleCops (or any other Web-security forum such as those listed in Chapter 13) and ask.



To check your installed programs online to see what they are and what they do, here are some Lists that can help you:

- Answers That Work at www.answersthatwork.com/home_page.htm. Click the Task List button. Although they say these programs usually appear in the Task Manager List, we have often used this resource for finding applications that appear in the Add or Remove Programs list.
- The Greatis Application Database at www.greatis.com/appdata/. This index can be used for Startup applications as well.
- The Bleeping Computer Uninstall Database at www.bleepingcomputer.com/uninstall/all.html. It has a search box in which you can enter the name any program that's visible in Add or Remove Programs to see whether it's essential, nonessential, or malware-related.

After you've removed the unused and nonessential programs from the Startup List and the Add or Remove Programs listing, don't be surprised if you find your computer booting and running faster and more efficiently. If it suddenly gets sluggish again, that may be a hint to give it the once-over in search of malware — because you've just made it more difficult for malware to hide. Unfortunately, rootkits won't become magically visible if all you do is inspect your files and folders. So the next sections get into cleaning up.

Using the Windows Disk Cleanup Utility

The Windows Disk Cleanup Utility is included in all Windows versions and cleans out downloaded program files (such as ActiveX controls and Java applets), temporary Internet files from Internet Explorer, discarded files hanging around in the Recycle Bin, and temporary files created by running and installing applications (those files can clutter up your system and slow it down without your even knowing they're in there). You can also automate the Disk Cleanup utility to delete these files at scheduled times and regular intervals. In this section, we show you how to perform both of these tasks.

Cleaning up your hard drive using the Disk Cleanup Utility

For Windows XP, follow this procedure for each hard drive on your computer (as listed in My Computer) to clean them up:

1. Click the Start button and choose All Programs ⇨ Accessories ⇨ System Tools ⇨ Disk Cleanup.



Alternatively you can open the utility by clicking the Start button, choosing My Computer, right-clicking Local Disk, and then choosing Properties. Then, on the General tab, click the Disk Cleanup button.

The Disk Cleanup window appears, asking you which drive it should clean.

- 2. If you only have one drive, click OK. Otherwise, select the drive you want to clean first, and then click OK.**

The next Disk Cleanup window that appears examines your hard drive to see how much space you can save by using the utility.

- The main Disk Cleanup window shows two tabs: Disk Cleanup or More Options.
- The Disk Cleanup tab shows you a list of items that can be safely deleted or compressed to provide more disk space. At the top you see the total possible space that can be made available. Click each item once to view its description. That info should help you decide whether you want to remove a particular file or folder.

- 3. Click an item in the list and then click the View Files button to see the files that can be deleted.**

You see a list of items with check boxes next to them; a check mark tells Disk Cleanup to delete those specific files after you click OK.

- 4. Put a check mark in the box next to each item you want to go away and then click OK.**

Disk Cleanup dutifully begins clearing space — and a progress window appears (this operation may take a few minutes). Disk Cleanup will close itself when it's finished.

Using the Disk Cleanup utility to clean up other things

When you open the Disk Cleanup Utility (as described in Step 1 and 2 in the preceding list), you'll notice that when the Disk Cleanup for C: (where C: is the drive you're cleaning) appears, there are two tabs on the window. What happens when you choose the Disk Cleanup tab is shown in the previous steps. Here's an opportunity to clean up even more hard drive space: the More Options tab, which helps you free more space by removing

- ✓ Unused Windows components
- ✓ Installed programs
- ✓ Old restore points left over from using System Restore

Removing unused Windows components

To remove unused Windows components from the More Options tab, follow these steps:

- 1. On the More Options tab, click the Cleanup button in the Windows Components section.**

The Windows Components Cleanup Wizard appears.

- 2. Click an item to view its description.**

If subcomponents are installed, the Details button is available; click this button to see a list of subcomponents — and if *those* have more subcomponents, you see the Details button again. Click it to see more.

- 3. If you know what each subcomponent is, and are certain that you don't need it, then click to clear the check box beside it; when you've unchecked all the stuff you want to get rid of, click OK.**

A progress box appears, followed by a box labeled Completing the Windows Components Wizard.

- 4. Click Finish.**

Removing installed programs

Clicking the Cleanup button in the Installed Programs section of the More Options window will (surprise!) open the Add or Remove Programs window. Use the Installed Programs option in Disk Cleanup to view unused programs. You can uninstall or modify their installation here as described in the “Removing unused programs” section earlier in this chapter.

Removing System Restore points with Disk Cleanup

Clicking the Cleanup button in the System Restore section of the More Options window summons a Disk Cleanup dialog box simply asking whether you'd like to delete all but the most recent restore point. Remember that System restore points are automatically cleared after 90 days. Even so, if you have restore points that you'd rather not keep, then click Yes. When you have fully cleaned your computer of any malware, especially if you've had it for some time, set a fresh restore point and do this again so that your most recent restore point is definitely trustworthy.

Defragmenting your hard drive

All Windows versions come with a built-in disk defragmenter. As you use your computer — creating and deleting files or installing and uninstalling

programs — gaps begin to appear on your disk. Windows fills those gaps by splitting some files into little chunks (depending on their sizes). The result is that files (and pieces of files) get scattered all over your hard drive — which makes the read or write heads work a lot harder to access individual files. The longer you use your computer without defragmenting the hard drive, the more fragmented it becomes — which has a hefty impact on your computer's performance and stability. To help you maintain an efficient computer, Microsoft has included the Disk Defragmenter utility in every version of Windows since Windows 95.

Disk Defragmenter works by moving large sections of data to another part of the disk in order to free up space, and then putting the fragments together as one contiguous piece. This process is called *defragmenting*. It continues this operation with each file until every file is assembled and put in its place.

Knowing what to keep in mind before defragmenting

Regardless of which version of Windows you're defragmenting, it's good to keep the following fine points in mind:

- ✔ If you have never defragmented your hard drive before, and you've been using the computer for a long time, be prepared for the procedure to take several hours, especially with hard drives of 20GB and larger. Let the utility do its job, and go do something else. Check back now and then to ensure that everything is okay. The more often you defragment, the less time it takes to complete.
- ✔ If you've turned off unnecessary applications with Windows Task Manager, you'll have to reboot or restart your computer when defragmentation is complete.
- ✔ Alternatively, you can restart your computer in Safe mode so the background applications won't be running. (We show you how in the "Speeding up defragmenting" section a bit later in this chapter.)
- ✔ Whenever you uninstall programs or clean up your machine, it's a good idea to defragment the hard drive afterward. The more you defragment the hard drive, the less time the process takes to complete.



Defragmenting Windows XP

If you're a Windows XP or 2000 user, you must be logged on as an Administrator (or be a member of an Administrator group) to use the defragmenter.

To defragment your Windows XP computer, follow these steps:

1. **Click the Start button and choose All Programs ⇨ Accessories ⇨ System Tools ⇨ Disk Defragmenter.**



Alternatively, you can click the Start button and Open My Computer, right-click a drive and choose Properties. Then click the Tools tab and (finally) click the Defragment Now button. You can also start Disk Defragmenter from the Run box. Click Start ⇨Run; then type **dfrg.msc** and click OK.

2. Windows XP Disk Defragmenter allows you to analyze the drive before you defrag it. Click the drive you want to check first; then click Analyze.

3. When the analysis done, check the results by clicking View Report.

The report tells you whether the drive needs defragmenting.

4. If you want to defragment your hard drive, click the Defragment button.

If this is your first defrag, then it's time to go load up the dishwasher or get the laundry done while you're waiting. Check on the progress now and then. After some time (especially the first time, which can take hours), defragmentation is complete. The graphical representation of your hard drive should look healthier than it did originally — with few or no red stripes visible on-screen.

5. When defragmentation is done, quit Disk Defragmenter, or defragment another drive.

6. When all defragmentation is complete, restart your computer.



For more information and help with defragmenting your hard drive please refer to the CastleCops Wiki on Disk Defragmentation at

http://wiki.castlecops.com/Disk_Defragmentation

Speeding up your defragmenting process

Defragmenting can be a long process, but you can speed it up somewhat. The following options can help you speed up your defragmenting process considerably:

✓ **Defragmenting in Safe mode.** Before you open the defragment program, restart your computer in Safe mode. Background programs, if they're running, can cause the Defragmenter to stop and start repeatedly — increasing the time it takes to complete the job. Using the Safe Mode method prevents that delay, and it makes it easier if you don't know which programs to turn off in Task Manager.

This method applies only to computers on which Windows XP is the only operating system.





To defragment in Safe mode, follow these steps:

- 1. Restart your computer.**
- 2. As your computer starts up, begin tapping the F8 key until the Windows Advanced Options Menu appears.**

Some models of computers may give you an error message when you do this. If you get such a message, restart the system and try again. If it doesn't work the first time, keep trying.
- 3. Use the arrow keys to select Safe mode from the menu (make sure you select it *without* network support) and press Enter.**
- 4. Log on using your normal user-profile name (if it has administrative privileges) or as the Administrator.**
- 5. Defragment your hard drive, as described in the previous section.**

To ensure optimal speed, don't do anything else on your computer while the Defragmenter is running.
- 6. When you're finished defragmenting, restart your computer and let it boot normally to return to Normal Mode.**

✓ **Temporarily disable your Internet connection.** If you have an "always-on" Internet connection, it's important that you be offline for this procedure. To temporarily disconnect your connection, follow these steps:

- 1. Click Start, choose Control Panel, and double-click Network Connections.**

Alternatively, in Category View, click Network and Internet Connections and then click Network Connections.

- 2. Right-click each connection and select Disable from the context menu.**

✓ **Disable your screen saver.** To prevent the screen saver from interrupting the defragmenting process, follow these steps:

- 1. Right-click the desktop and select Properties.**
- 2. Select the Screen Saver tab. Set the screen saver to [none] in the selection box.**
- 3. Click Apply and then click OK.**



Checking out alternative defragmenting utilities

The Windows XP Disk Defragmenter is really a Lite (or scaled-down) version of Diskeeper by Executive Software. It does the job well enough, but if you prefer a defragment program with all the bells and whistles, many third-party programs are available. Two of them are

- ✓ **Diskeeper Pro.** The authors are both using the full version of Diskeeper Pro. Diskeeper Pro allows defragmentation to be run as a scheduled task, and can also be used to defragment networked computers remotely.
- ✓ **PageDefrag.** Disk Defragmenter does not defrag the page file, Registry hives, event logs, nor hibernation files (usually on laptops and notebooks). Most third-party defragmenters won't do them either. Sysinternals provides a free application that will defrag all of them called PageDefrag. It's compatible with Windows NT 4.0, Windows 2000, Windows XP, and Windows Server 2003. Here's a link that offers more information and downloads:

www.sysinternals.com/Utilities/PageDefrag.html

Using Registry cleaners

The Registry is the heart and soul of your Windows computer. It does for your files and software what your CPU does for your hardware. The Registry is one of the most complicated and sensitive groups of files in your operating system; you need to be an expert to edit it manually. An advanced user may feel confident with using one of Windows Registry tools (such as Regedit or Regedit32), but the rest of us mortals have to watch our steps pretty carefully. Be cautious in your choices of commercial utilities for editing the Registry — some are real grouches, neither user-friendly nor system-friendly.



Do not attempt to edit the Registry manually unless you've backed it up beforehand and are confident that you know what you're doing. As stated earlier, errors of any kind are not forgiven there. Examples include a missed comma or a space added between characters can result in disaster for your computer's stability. Everything, including spaces and characters have a specific meaning in the Registry. See the "Backing up your Registry" section earlier in this chapter for the lowdown on how.

Always use an application to make edits to the Registry, such as removing old file remnants from uninstalled programs, redundant paths, and obsolete keys.

Check Bonus Chapter 2 for several registry cleaning utilities we recommend, along with instructions on how to use them.

Controlling Removable Devices

Part of good computer hygiene is making sure your computer is difficult enough to corrupt that someone can't simply walk up, pop a CD into your drive while you're not looking, and walk away having started an exploit without even having to run the CD. In that scenario, Windows AutoRun feature would start the CD up and infect your computer automatically — nobody would have to do anything besides put the disc in. The same goes for the AutoPlay feature, which automatically opens a new disc, USB device, or other removable medium as soon as one is inserted into your computer. The following sections show you how to protect Windows by disabling these dangerous (if convenient) features.

Disabling AutoRun

Sometimes the convenience of automation just isn't worth the security risk. Here's a classic example: When AutoRun is enabled on your CD or DVD drive, a setup file called AutoRun.inf automatically loads — and instantly runs — the software on the disc. That means any lurking malware or DRM (Digital Rights Management) tools can go into action from the get-go. (This is exactly how the Sony rootkit worked.)



TIP

For all versions of Windows, you can just hold down the left Shift key to turn off AutoRun as you insert a CD or DVD in the drive. However, it's better (and less of a pain than having to remember the Shift key all the time) if you disable AutoRun completely. That way, if some stranger wants to load stuff on your laptop or desktop PC without your permission, it's no dice. For Windows XP SP2 Professional, follow these steps to disable AutoRun completely:

1. **Right-click the icon for one of your CD and DVD drives and choose Properties.**



REMEMBER

You have to perform these steps for each CD and DVD drive you have on your system. Otherwise AutoRun does its thing on any drive you haven't told *not* to use it.

2. Click the AutoPlay tab.
3. Select each type of CD or DVD one at a time (Music CD, DVD, Mixed Content, etc.), click Select an Action to Perform, click Take No Action, and then click Apply.

Do this for each type of media on which you want to disable AutoRun.

4. When you're done, click OK.

Turning off AutoPlay on all external drives and devices

If you take your computer to public places often or travel with it, you may want to disable AutoPlay on external drives and devices to improve security. This section shows you how to turn off AutoPlay on all external drives and devices at once.



Online you can find various Registry tweaks that can disable AutoPlay on all your peripheral and optical drives — including USB and floppy drives — but be cautious with those tweaks. Always have a backup of your Registry and a new restore point ready *before* you apply them. Registry changes, if not done precisely right, can render your computer useless.

On Windows XP Pro SP2 you can turn off AutoPlay on all your external devices at the same time without tweaking the Registry. To do so, you need to be an Administrator or a member of that group to follow these steps:

1. Click the Start button and click Run.
2. Type `gpedit.msc` into the Open: field and click OK or press Enter.
3. Choose Computer Configuration ⇄ Administrative Templates ⇄ System.

The System level becomes accessible, offering several options. Scroll down on the right pane.

4. Enable the Turn off AutoPlay option by clicking it. Click it once to view its description on the left side of the right pane. To turn it off, double-click it.

Note that this does not turn off AutoPlay for music CDs. You can disable programs from loading from external media, but still enjoy your tunes. If you have multiple accounts on your computer, the settings are changed for all of them.



Living without AutoRun

Okay, life is a little less convenient without AutoRun. Now, whenever you want to play a CD or DVD, use your media-player software to open the desired files on the disk. For playing game discs, you'll have to open My Computer on the Start menu locate the game's main file folder manually. But you can get back some of that convenience without allowing AutoRun be a klutzy gatekeeper. Just create your own shortcuts and move them to the desktop. Here's the way to do that, using a game's main file folder as an example:

1. Open My Computer on the Start menu and browse to the game's program folder.
2. Find the primary executable file for that game, right-click it, and choose "Create New Shortcut" from the context menu.
3. Click and drag on the shortcut to your desktop.

After you're done, because most games require their CDs for you to play them, you'll simply put the game CD in the drive, and then double-click the game shortcut on your desktop to start your game.

Scanning boot sectors before using external media

Going along with the previous section, when you insert a new CD or any external media into your computer, you're playing with fire if you don't scan it first. Your computer is most vulnerable at bootup or startup — this is when boot-sector viruses can appear and spread. You can get them not only from Internet downloads, but from anything that fits into a drive on your computer — CDs, DVDs, USB storage devices, and floppy disks. That's because most removable media have either a *boot sector* or *setup file* — the first place the computer looks to get instructions on what to do with the stuff on the medium. A black-hat hacker just couldn't ask for a more insidious place to hide a rootkit or other malware.



Boot viruses that hide in the boot sector can render a hard drive unusable without even breathing hard. That's why it's so very important to scan all downloads and removable media with your antivirus and other anti-malware applications, *before* you open or run anything. Turning off AutoPlay makes this easier.

Note that many antivirus applications scan the boot sectors of all present media at startup, but it's still important to scan the rest of the files on the disk, just in case.

Chapter 4

Staying Secure Online

In This Chapter

- ▶ Stopping JavaScript and ActiveX scripts in their tracks
 - ▶ Developing savvy and secure Internet practices
 - ▶ Getting your browser on your side
 - ▶ Erecting a firewall
 - ▶ Protecting your system with scanners
-

The Internet gives you and your computer many points of contact with other people and machines, far more than you have in the physical world of tangible things — all of it happening at the speed of light. But any contact that you and your computer engage in online has the potential for harm to your computer — and even to you, your family, and business. Besides the harmful people online, many automated dangers exist such as viruses, worms, and other forms of malware. The Internet is the number-one route for malware to find its way to your computer. This chapter shows you how to fortify your operating system so you can enjoy the Internet, and still keep your computer malware-free.

Good Practices Are a Good Start

How do people who use the Internet get malware or rootkits onto their computers? The easiest way to get malware is to go online without any protection, simply assuming that the Internet is a friendly place, all about networking and meeting new people. The best victims rarely get patches and updates but download *lots* of programs to help search for and visit all those wonderful Web sites, let everyone they meet know how to contact them via e-mail or IM, respond to all, and open every attachment. They always fill out all the forms they find, use easy passwords (or no passwords at all), enter every contest — and then wonder why they get infected.

Bottom line: Getting a rootkit usually involves some user action. For example, a rootkit can be used to administer Digital Rights Management (DRM) software on a music or videodisc (such as the Sony rootkit), as has been illustrated in the news. The users simply play the disc on their computers — and the DRM rootkit is installed automatically. The same thing can happen if an infected e-mail attachment is carelessly opened. That attachment may contain a trojan that has been tailored to deliver a rootkit.

For as many threats as the Internet harbors, preventing yourself from becoming a victim is totally under your control, and it simply involves exercising common sense. The following sections give you an overview of guidelines to live by when surfing the Internet; they can help keep you safe online.

Choosing your contacts carefully

What constitutes a “contact” online? For the purposes of this chapter, a *contact* is any type of connection to your computer. A few definitions may help. “A list of people in various name and address programs” is the noun form; as a verb, contact means “to be in or establish communication with,” and in technical terms (again as a noun) it’s “a junction where two electrical conductors connect.” Every one of those meanings is relevant here. Online you have contacts whenever you log on to a network or your Internet connection, surf with your Web browser, use your e-mail, visit a Web site, use an Instant Messaging (IM) program, use Internet Telephone or Voice Over Internet Protocol (VoIP), participate in videoconferences, use Multimedia Message wireless services (MMS) with your personal digital assistant (PDA), or download programs.

Contacts online are far more numerous than those you share with other people. Every time your computer links or interfaces with other computers, you have contacts. The Internet is a huge number of interconnected computer networks; they’re like cities and towns linked by highways and roads. If somebody drives over to your house, assuming you know each other, then you have contact. What happens if you *don’t* know each other? You may still have contact — provided your visitor has identification that shows that he or she is entitled to be there. When it comes to protecting yourself from harmful contacts, it’s important to realize that what you connect with is as essential as with whom you connect.

The point is just basic common sense: Choose your contacts carefully. Giving out your e-mail address or IM to the wrong person or connecting to the wrong server by visiting a malicious Web site can result in spam, malware infections, and many other problems.

Getting trojaned online

When we ask you to take due care and caution while surfing online, we know because we've been there already. One day, some years back, Larry got it into his head that he wanted a specific old-style terminal-emulation program. He discovered that it had become *abandonware* (software that's no longer supported) and it wasn't available in North America. So he searched everywhere else. He found a link on it to a Web site on the Asian continent. The site was partly in English, so he was able to navigate. He found a link that offered more information about the emulator. He clicked it. Suddenly CPU usage shot up to 100 percent and the browser froze. Using a keyboard shortcut, he turned the browser off. The screen was black, but the computer was still running, and the drive was going bonkers. Scary!

What would you do? Larry turned off his modem first, and then turned off the computer manually. Instant hard-drive crash. He waited for several

minutes to catch his breath before turning the power back on for his computer. He left the modem off and restarted in diagnostic mode (Safe mode — see Chapter 3 to see how to do that). He began running all of his scanners, starting with the antivirus. The anti-trojan scanner (TrojanHunter) found the problem: Simply by clicking that link, he'd got himself a dandy little backdoor trojan and a keylogger (more about those in a minute). Some system files were damaged as well, as discerned by System File Checker. He was able to remove the trojans and fix the files, but he stopped doing searches at Asian Web sites thereafter. Once burned, lesson learned.

Why should you have to go through all that hassle? A lot of people in the security communities got there by going to the school of hard knocks. You don't have to go to that school if you don't want to. We can help you to help yourselves. It's the least we can do.

Surfing safely

Learning to start surfing safely and securely begins with a properly configured browser. You can discover what to do about that in the later section "Bashing Your Browser into Submission." Yet, the guidelines, tips, and instructions we provide are not enough by themselves to ensure your safe excursions on the Internet. It all starts with you. No matter how experienced and savvy a computer user you are, eventually you'll encounter something (ahem) *unexpected* while online.

The following list gives you some general guidelines to keep in mind when surfing the Internet, and the following sections discuss guidelines specific to downloading, e-mail (and attachments), and instant messaging.

- ✔ **Don't leave your computer online constantly.** As we've mentioned before, how desirable a target you make yourself is largely in your own hands. An always-on broadband connection makes you a very desirable target. If you have a broadband or DSL Internet connection, don't leave your Internet access turned on or connected all the time; shut down your system (or remove the Internet cable) when it's not in use. Leave your cable or DSL modem or router on so you don't have problems re-connecting to the service. Disable your Network Connections via Control Panel, then unplug the network cable at the back of your computer. This method is described in detail in Chapter 3.
- ✔ **Whenever possible, conduct your online activities using a Limited User account.** Because rootkits need administrative access to inflict their harm, surfing as a humble Limited User offers an extra level of built-in protection. Some rootkits can contain code that allows them to bypass restrictions imposed by group policy settings, or even upgrade privileges. Still, surfing as a Limited User presents an added obstacle that makes you a less desirable target from a hacker's perspective. The "Establishing limited-access user accounts" section later in this chapter explains how.
- ✔ **Use Encrypting File System (EFS) to keep your private files safer.** EFS is available by default and only on NTFS formatted systems, Windows 2000, XP and higher. You can encrypt any files and folders except the system ones. To protect your private files while online, open My Computer then look at the "Files Stored on this Computer" at the top. Find the folder corresponding to your user account, right-click it and select Properties. Click the Advanced button. Put a check in the box beside "Encrypt contents to secure data," and click OK. You will be asked if you want to include sub-folders and files. Yes, you do. You will still be able to use your files and folders as long as you are logged in to the user account in which they were encrypted. When you open one, it's decrypted automatically, and then re-encrypted when you close it. These files and folders are not available to anyone else. As far as they're concerned, you're in the john.
- ✔ **Surf warily if you're using public computers.** In the case of a system you have no control over — and can't log on to as an administrator — you should *never* log on to anything of importance (bank accounts, credit card accounts, corporate e-mail, and so on). In case of an emergency where you might have to check an e-mail from such a system, you should immediately change the passwords on those accounts as soon as you return to a safe system.
- ✔ **Stay away from pornographic Web sites.** Porn Web sites are some of the most flagrant distributors of malware. Beware of sites that invite you to download free content or install something that allows you to view pornographic content. Remember, nothing is free! Porn sites are infamous for being one of the most common sources of infection.

Downloading safely

One of the easiest ways to reduce your probability of infection is to avoid downloading any new applications. Realizing that total isolation isn't the most practical approach, we've assembled the following list of guidelines. They'll reduce your chances of downloading a rootkit — or any type of threat — unintentionally:



- ✔ **Don't use peer-to-peer (P2P) networks.** No, we're not just being killjoys. Peer-to-peer (P2P) networking (the original Napster and Kazaa being two of the more popular examples) can introduce trojans or worms into your P2P download folder that can literally unleash a torrent of threats onto your system. Using P2P file-sharing is an invitation to disaster and is illegal. P2P file-sharing sets your computer to act as a server to which *anonymous others on the network have unlimited access*. Plus, anything you download could contain a disguised trojan or a rootkit. Similarly, stay away from pirated sites that deal in illegal distribution of copyrighted material or cracked software, such as Warez distribution sites.

A known rootkit worm called W32.Fanbot.A uses advanced DKOM techniques (discussed in Chapter 7) to hide successfully on Windows 2K and XP systems. P2P networking is one of the known major vectors by which it spreads.

Be very frugal when it comes to accepting ActiveX controls. ActiveX controls represent downloaded program files. Malware distributors will often use them to try to con you into downloading their crapware. The Gromozon rootkit very cleverly prompts Internet Explorer users for approval to download a file from a famous search engine: `www . google . com`. (Note the spaces in the bogus address.) Naturally, most users impulsively agree to download something that sounds so familiar — yet, unbeknownst to them, they will have agreed to install the latest cleverly disguised adware rootkit.

Don't put anything on your system without thoroughly checking it first. Before you download anything, perform a Google search on the product's name to make sure reputable sources consider it trustworthy. The Internet is anonymous; anyone can try to lure you in with a glitzy Web site. A polished-looking site is no guarantee of a trustworthy product; the site may be engineered to lure you into downloading infected material.

- ✔ **Read all End User License Agreements (EULAs) and Privacy Policies before downloading anything.** You may not realize what you're getting yourself into unless you do that. Did you know (for example) that some EULAs make *removing their software* a violation of their license agreement? Talk about one-sided! Even installing programs from CD can pose a threat. Nancy recently installed a mouse driver from a vendor-supplied CD — and found that a recognized adware threat came bundled along with it. Luckily, the active-protection component of her anti-spyware program caught the intruding program right away.

E-mail and attachments

As you know, e-mail is nothing like the letters you get from the postal service. E-mail is like a postcard; anyone can read it while it's in transit. The return addresses can be *spoofed*, or altered to hide the identity of the person who sent it. Attachments are files added on to e-mails and can be deceptive — Grandma's chocolate chip cookie recipe could also contain unexpected "ingredients." Below are some tips to help you enjoy e-mail safely, with greater security.

- ✔ **Use strong passwords for your e-mail accounts.** All passwords also need to be changed regularly, at least once a month. E-mail accounts without strong and changing passwords can be more easily hijacked and used for bad purposes. For tips on creating well-muscled passwords, see the "Developing strong passwords" section later in this chapter.
- ✔ **Delete unknown e-mails immediately.** If you get an e-mail from someone you don't know and aren't expecting, send it to the trash unopened. Larry has received e-mails from "Mom," but his real mother doesn't have a computer and doesn't use the Internet. Those always go in the Recycle Bin right away. (Oh, yeah . . . don't forget to empty it.)
- ✔ **Don't use a preview pane.** If your e-mail program has a preview pane, then disable it so you see only the titles of the e-mails — and they remain closed unless you deliberately open them. Preview panes open all e-mails first, so you can see the first few sentences. Once they're opened, it's too late; if any of them is infected, so is your computer. The titles of unopened e-mails cannot infect you.
- ✔ **Be wary of any e-mail that purports to be official or of some dire emergency nature.** People — especially those who have seriously important business with you — will contact you using a more trustworthy and valid method than e-mail if a *real* emergency is going on. Check with your banks and credit-card companies to learn their policies toward e-mail. You can arrange with them to be contacted by registered postal mail for any important issues. Then, if you get e-mails that appear to be from your bank or credit card companies, you can delete them unopened.
- ✔ **Don't download anonymous e-mail attachments.** Better yet, go one step further — don't download *any* e-mail attachments, even from known sources, unless you've been notified to expect an attachment beforehand. If you must download an attachment, save it to a folder first. Don't open it. Find it, right-click it, and choose one of the anti-malware applications listed in the context menu. Scan the attachment with each one until you're satisfied it's safe.
- ✔ **Don't even trust e-mails from people you know.** If a person you know and trust sends you an attachment, scan it with antivirus, anti-trojan, and anti-spyware applications before opening it whether you're expecting it



or not. Most anti-malware applications provide the means to scan single files, but you can forego that scanning step if your e-mail client — or your anti-malware application — scans e-mail attachments automatically before they're downloaded.

Just because you know an e-mail's sender doesn't necessarily mean it's safe. Malware purveyors can easily steal an e-mail contact list and then send infected e-mails using a familiar name that you recognize. The bad guys can even steal message-header information so the message subject matches one you've used in previous exchanges. Of course, all these gimmicks are aimed at persuading the recipient to open the e-mail without scrutinizing it first. Nancy recently received a series of infected e-mails that were labeled with stolen subject headings she had used before in her correspondence. Analysis revealed that every one of the infected e-mail attachments contained a copy of the KamaSutra (aka MyWife) e-mail worm.

- ✔ **Be cautious to whom you give your e-mail addresses.** People who have your address on their computers could become infected with a mass-mailing worm and send it to you. Use one or more free Web e-mail clients for the places you need to provide addresses. Many of these provide free anti-malware scanning. Keep your private, paid-for e-mail addresses for the select few.
- ✔ **Avoid contests online.** Contests are often used as an excuse to obtain e-mail addresses so somebody will have a place to send advertising or spam. Additionally, this holds true for “free” offers — or anything that requires your e-mail address without good reason. Free Web e-mail comes in handy here. If an address gets inundated repeatedly with useless spam, close the account and open another.
- ✔ **Scan files before opening them.** No harm comes from simply downloading or saving files. The trouble starts only when you *open* them; opening an attachment can be the same as starting a program. Most e-mail clients offer a scanning feature that will relieve you of the responsibility of manually scanning your e-mail attachments. Many antivirus products also offer e-mail plug-ins that automatically scan e-mails for threats. However, if *neither* your e-mail client nor your antivirus provides this built-in layer of protection, you should follow this procedure when receiving e-mail attachments that you're expecting:
 1. **Download the attachment by clicking it and saving the file on your computer.**

Do not open the file yet.
 2. **In Windows Explorer, right-click the file, choose Scan Selected Files with . . . , and specify your anti-malware scanner.**

When the scan is complete, your scanner will let you know if the file is clean. Scans take only a moment but can save later frustration and hardship.

Alternatives to this procedure exist where the files in question are scanned on a server separate from your computer before you download them. Note that no scanner is perfect. An outside chance always exists that your attachment may contain unknown malware.

Instant messaging

The advice for dealing with e-mails (discussed in the previous section) goes double for instant-messaging applications such as IRC, ICQ, MSN, AIM, Yahoo! Messenger, and so on. Keep the following things in mind when using IM:

- ✔ **Be careful with whom you share your handles or nicknames.** Many teenagers pride themselves on having a hundred or more contacts in their Buddy Lists as a matter of social prestige. But it's also a recipe for malware disaster, or worse.
- ✔ **Be wary of using IM file-sharing features.** Most IM applications allow peer-to-peer (P2P) file shares, which can allow unscrupulous people easy access to your files. Disable file sharing in your IM applications to prevent malware compromise and hacking attempts by this route.
- ✔ **Be sure to have your anti-malware Real Time Monitoring scanners on whenever you use IM.** Doing so prevents malware exploits during your chat sessions. These can come via scripts introduced by chat partners, and links to malware or compromised Web sites.
- ✔ **Use a valid auditing policy.** Chapter 6 and Chapter 8 explain the ins and outs of auditing. The idea is to monitor critical areas of your computer to prevent hacking attempts and malware exploits. Auditing can also help you if you're backtracking through the records to discover the source of malware and hacking attempts.
- ✔ **Don't click links embedded within instant-messaging text — even if they're from your "buddies" — unless you're 100 percent sure your buddy meant to send it.** Most IM networks have been targeted to spread malware; the accompanying sidebar shows how that works. Because IM is such a common route of infection, you should — at minimum — configure all your instant-messaging programs to require approval for any incoming message from an unknown "buddy."
- ✔ **Use a chat client to monitor and maintain all your IM accounts.** For example, Trillian (by Cerulean Studios) has many security features that individual IM applications lack — such as enabling encrypted chat sessions, chat rooms, multiple simultaneous sessions, audio chat, and more. It has both freeware and for pay editions. You can check it out at the following URL:

www.ceruleanstudios.com/learn



With “buddies” like these . . .

The AOL Instant Messaging network was recently used to propagate a widespread rootkit worm attack dubbed by some as the AIM virus. Worst of all, this particular worm stole the buddy lists of AIM members and used their names to send the infected messages. Victims became

infected when they clicked innocent-looking embedded URLs within the instant messages — which they assumed were intentionally sent from their buddies. The AIM virus propagators used clever (though sinister) social-engineering tricks to dupe AIM users.

Developing strong passwords

We cannot emphasize enough how important well-chosen passwords are to your security. In this section, we describe how to craft good passwords.



A great little program called Any Password is on the *Rootkits For Dummies* DART CD for you to use. It can help you to create, store, and maintain your password files in encrypted folders using Master passwords that you select. For those who are loathe to use their imaginations, software is also on the CD that can help you generate passwords of any length and complexity. See the Appendix and Bonus Chapter 2 for more information.

Strong passwords are easy to make if you use a little imagination. Try using a phrase from a favorite nursery rhyme, such as “Mary had a little lamb!” Now we add a number to it — substituting 1_lamb for “a lamb,” we get `Mary_had_1_little_lamb!` That’s a 23-character password if we use the underscores for spaces (and we *should* use them if underscores are allowed); it’s 19 characters if we don’t use ‘em. Here’s what the results look like:

```
Mary_had_1_little_lamb!
```

```
Maryhad1littlelamb!
```

This makes for a good, strong password that’s easy to remember. If you wanted to make it even longer and stronger, you could add something funny or silly at the end — say, Bam Bam! We end up with `Mary_had_1_little_lamb!_Bam_Bam!` Now we have a 32-character password.

For those of you who are interested in creating your own passwords, the following list gives you some more tips and things to consider when making them.

- Don’t use words you find in dictionaries. Known words can be easily and quickly searched for and compromised. (If you mix them up with numbers and symbols as we did here, that works.)

- ✔ Don't use personal information in your passwords. No birthdates, pet names, family names, addresses, and so on.
- ✔ Use symbols and special characters such as ! @ # \$ &. Just don't use them to spell out cuss words (that's one of the easiest types of password to crack).
- ✔ Use pass-phrases instead of passwords. *Danger to health increases the more you smoke and inhale* is a good example, and so we get D2HiTmUS&I — which is also more than eight characters. Very hard to crack. (And no, we don't publish our real passwords in our books — or anywhere. But you knew that, right?)
- ✔ Using a nursery rhyme or the title to a favorite song, using variations for different passwords, will help you remember them should you forget later on. Most important is never using the same password for different secure settings (for example, your bank account *and* e-mail account).
- ✔ Using Any Password (as described in Bonus Chapter 2 and the Appendix), you need only remember one or a few Master passwords to recall them all. You can use copy-and-paste from it when inputting long passwords where you need them. Any Password automatically clears your Clipboard when you close it.
- ✔ If you're not feeling very creative (or would like to leave the job to someone else), there's an excellent online random-password generator at www.randpass.com.



For more techniques on developing effective passwords and password policies, check out *Perfect Passwords* by Mark Burnett (Syngress Publishing).

Establishing limited-access user accounts

In Windows XP Professional, you have a variety of user accounts available to you. The Administrator account can do everything, and can make changes to all the other accounts, such as for Power User or Limited Access. On the other hand, malware normally picks up the access capabilities and privileges *of the account you're using at the time it's introduced*. So if you're surfing online with an Administrator account when your computer becomes infected, the malware gets unlimited access to your system.

To reduce and prevent damage to your system from malware, it's a good idea to establish limited-access user accounts for all surfing online. Limited-access accounts have the least privileges of all the accounts available to you. Using a limited account prevents many malware applications from being installed on your computer; anything that does make it onto your system will be severely restricted in its effects. Of course, you'll still have to log on as an administrator

for essential tasks such as installing new applications, downloading Windows updates, and running malware fixes (to name a few), but it isn't wise to be logged on as administrator all the time.

To create a limited user account, follow these steps:

1. While logged on as an administrator, click the Start button, click Control Panel, and open User Accounts.

Alternatively, you can open My Computer, click Control Panel and open User Accounts.

2. Click Create a New Account.

You see a place to type in a new account name. Try to use a name that's not descriptive of the purpose of the account.

3. Type in a name and click Next.

4. Select the Limited account option and then click Create Account.

You see a newly created account listed alongside the existing accounts.

5. Click the new account.

The next order of business is to create a password for the new account.

6. Click Change the Password in the list.

Ignore the warning given at the top. It applies to established accounts with existing passwords. We are making a first password for our new account. By clicking the headings under Learn About (on the left side), you can read information and tips on creating secure passwords.

7. Type in a strong password in the first box; repeat it exactly in the second box.

8. If you want a password hint, type it inside the last box. If you don't want one, leave it blank.

9. When you're ready, click the Change Password button.

You're returned to the Account Changes screen. The User Accounts wizards are easy to use. You can go back, cancel, or change anything about this account until you're satisfied with the results — choose a new icon for it, change the name or password, delete it, whatever. When you're done, you can make changes to a different account — or create a new one under Related Tasks on the left.



While you're at it, disable the Guest account (which allows anyone to gain access to your computer). From the main User Accounts screen, click the heading for Guest in User Accounts and select the Turn off the Guest Account option.

Using a HOSTS file

HOSTS is the name of a file without an extension that's included with every Windows operating system. When you click a link to a Web site, your computer first checks your HOSTS file for the IP (Internet Protocol) number for that address. If the HOSTS file doesn't have that IP, then the computer asks the DNS (Domain Name Server) via your ISP (Internet service provider) to map the Web site name to an IP address. (For example, one of the server addresses to CASTLECOPS.COM is 209.213.221.100.) Your computer first asks the HOSTS to provide the IP number for the alphabetical address. If it does not get it, then it asks the DNS instead. The idea is if you put the addresses of the sites you visit most frequently in your HOSTS file, it will speed up your surfing by eliminating the IP-to-domain-mapping step normally done by your DNS.

The HOSTS file was first conceived of as a way to hasten Web surfing, but with the advent of malware, another beneficial use soon became apparent: The content of your HOSTS file can be modified so it actually protects you while you're surfing. Here's how that's done: Addresses to known malicious or illegal Web sites are inserted into the HOSTS file and assigned an IP number of 127.0.0.1, which is the address of your local computer. Any attempt to navigate to one of these unsavory sites will result in an immediate redirect to your own computer. This enables a carefully tailored HOSTS file to protect you by automatically blocking access to suspect Web sites.

A *blocking HOSTS file* will relieve *you* of much of the responsibility of distinguishing good sites from bad; it makes that decision for you (provided you've already listed those sites in the HOSTS file). A number of well-maintained blocking HOSTS files are available for free download. One award-winning MVPS HOSTS file — one of the most comprehensive — can be downloaded here:

www.mvps.org/winhelp2002/hosts.htm



Spybot Search & Destroy on your DART CD has an immunize function that will prevent your HOSTS file from being altered by malware, if the Lock Hosts File Read-Only as Protection Against Hijackers option is checked. Spybot also offers its own version of a HOSTS file list under its Advanced menu options. Further instructions about using Spybot S&D can be found in Bonus Chapter 1 and in the Appendix.

It's best to use a tool to edit your HOSTS file, as it cannot be edited like a text file. To insert the Spybot Search & Destroy Hosts list into your HOSTS file, follow these steps:

- 1. From the main Spybot interface, choose Mode ⇄ Advanced.**

A warning appears, informing you about the dangers of making this change.

2. Click Yes.

Your main window menu is now in Advanced mode.

3. On the left pane click Tools.

A list of tools used by Spybot appears on the right pane.

4. Find and put a check mark in the box beside Hosts File.**5. Click the Hosts File icon that appears in the left pane.**

Everything listed in your HOSTS file appears in the right pane. Malicious Web site names are on the left side with the address to your own computer (127.0.0.1) on the right side. At the top you will see some choices you can make.

6. To insert the Spybot Hosts file list into your HOSTS file, click on the green + sign beside Add Spybot-S&D hosts list.

It's as simple as that.

You can also remove it if you wish, or you can edit your HOSTS file by clicking on an entry and selecting "Remove selected entries." If you make a mistake, you can use "Restore backup" at the top to correct it.



Bashing Your Browser into Submission

If your browser acts like a naive butler who lets just anybody in, then it's the major conduit through which junk and malware files arrive on your computer. To prevent junk and malware from piling up while you're surfing the Internet, you have to secure your browser — especially if it's Internet Explorer 6, the most popular malware target. Even if you use another browser, such as Mozilla Firefox, or Opera, you still have to use IE in certain cases (such as for Windows Update or installing software that requires legitimate ActiveX controls). The following sections show you how to configure Internet Explorer to limit your exposure to dangerous stuff.

Saying no to Java, JavaScript, and ActiveX

Java, JavaScript, and ActiveX are software systems that allow computer programs to be combined with Web pages, making Web site programs interactive and more enjoyable. However, these programs present inherent security risks to unwary users; they can act as vectors for viruses, trojans, and rootkits, ushering them into your computer without a peep (depending upon how lenient your browser settings are).

When you browse to a Web site without fully securing Internet Explorer, you may encounter pop-up windows that ask whether you'd like to download a particular component for use at that site. Most people simply click Yes or OK without realizing the implications of what they're doing — giving the unknown ActiveX or Java component complete permission to do almost anything to their computers. (Does that make you queasy, too?)

To help secure your computer against these threats, in this section we explain how to adjust the content zones in Internet Explorer 6. First, we show you how to adjust the settings for the Internet zone; then we guide you through some changes to the Trusted zone so you can add sites that you use often (and that you're absolutely confident don't pose a threat).

Follow these steps to optimize your security with the Content Zones in Internet Explorer 6:

1. Start Internet Explorer and choose Tools ⇨ Internet Options.

The Internet Options dialog box appears.

2. Select the Security tab.

Internet Explorer lets you set different levels of security for four different types of sites: Internet, Local Intranet, Trusted Sites, and Restricted Sites.

3. Select the Internet zone.

4. To make the Internet zone more secure, click the Custom Level button, and change settings as outlined in Table 4-1.

By applying these settings, you can automatically refuse Java and ActiveX pop-ups at every Web site you visit. JavaScript is included in the "Active Scripting" setting (see Table 4-1).

Table 4-1 Internet Zone Settings that Block Java, JavaScript and ActiveX

<i>For This Setting</i>	<i>Choose This Option</i>	<i>What the Setting Does</i>
Download signed ActiveX controls	Disable	Blocks signed ActiveX controls
Run ActiveX controls and plug-ins	Disable	Blocks ActiveX controls and plug-ins
Script ActiveX controls marked safe for scripting	Disable	Blocks scripted ActiveX controls
Font Download	Disable	Prevents hidden scripts

<i>For This Setting</i>	<i>Choose This Option</i>	<i>What the Setting Does</i>
Java permissions	Disable Java	Blocks Java
Allow META REFRESH	Disable	Blocks redirects to other Web pages
Display mixed content	Disable	Blocks unsecured content on Web pages
Drag and drop or copy and paste files	Disable	Prevents drag, drop, copy, and paste operations for all files
Installation of desktop items	Disable	Blocks installation of desktop items
Launching programs and files in an IFRAME	Disable	Prevents hidden IFRAME files and programs from running
Navigate sub-frames across different domains	Disable	Prevents browser hijackings
Software channel permissions	High Safety	Gives high-security settings to software channels
User data persistence	Disable	Disables save and load operations
Active scripting	Disable	Blocks JavaScript; prevents viruses and worms
Allow paste operations via script	Disable	Blocks Web sites that obtain the contents of your Clipboard
Scripting of Java applets	Disable	Blocks scripts in HTML Web pages
User Authentication	Automatic logon with current username and password	For Web sites using IIS, Internet Information Services only



You may want to allow *meta-refresh* — a form of automatic refreshing commonly used by news sites, forums, and legitimate download sites.



Consistent with the premise that you're severely locking down IE, it may be a good idea to change the User Authentication setting from Log On to Prompt for User Name and Password. This handy trick will increase your system's security.



Changing your settings as outlined in the previous list affects the Web sites you visit — a frequent result is that some Web sites won't work at all and others won't work very well. But hang in there — don't give up on changing your settings, and don't allow active scripting on every site by default. If you want to be able to run Java, ActiveX, or Scripting on certain Web sites, just add them to the Trusted sites zone (as discussed in the next section).

Adding sites to your Trusted zone

Adding sites to your Trusted zone enables you to browse sites you know are trustworthy without making continuous and repetitive decisions about what components and controls to trust on particular Web sites. When a site is in your Trusted zone, you can choose to enable Java and ActiveX controls and plug-ins for that site only, even though you're not enabling these controls for sites on the remainder of the Internet. Those you can pick and choose (as described in the previous section).

To add Web sites to the Trusted zone, follow these steps:

- 1. Open Internet Explorer and choose Tools ⇨ Internet Options.**
- 2. Select the Security tab and click the Trusted Sites icon.**
- 3. Click the Sites button.**



By default, you can only add secure sites here (sites using `https://`). To change that, in the Trusted Sites dialog box that appears, just uncheck the *Require server verification* (which specifies that `https:` isn't required) for all sites in this zone, and you can add any site.

- 4. Type (or paste) the URL for a site you want to add into the Add This Web Site to the Zone box and then click the Add button.**

You can add multiple sites if you want.

- 5. When you're done adding sites, click OK in the Trusted Sites dialog box — and then do so again in the Internet Options dialog box.**

Securing Opera and Firefox

Since the Opera and FireFox browsers don't support ActiveX, you can more easily secure them simply by adjusting their settings to block Java and JavaScript by doing the following:

For Firefox, with the browser open, choose Tools ⇨ Options ⇨ Content and then uncheck Enable JavaScript and Enable Java.



To further protect you, the NoScript Firefox extension allows JavaScript, Java, and other executable content only for the trusted domains you specify. Using NoScript protects you from *zero-day exploits* (vulnerabilities exploited the day they're publicly announced). The NoScript add-on may be downloaded here:

<http://addons.mozilla.org/firefox/722/>

Disabling JavaScript and Java within the Opera Browser is pretty quick: With the browser open, choose Tools ⇨ Quick Preferences and then uncheck the Enable JavaScript and Enable Java options in the pull-down Preference menu.

Disable AutoComplete in Internet Explorer

Internet Explorer has a handy-but-dangerous feature that automatically completes usernames and passwords on sites that you visit. To be safe, you should disable this feature. AutoComplete may seem convenient, but it can easily give away identifiable information. Anyone who logs on to your computer can access your passwords and open sensitive information.



You can use Any Password to copy and paste your passwords from lists you keep encrypted on your desktop. Any Password stores your passwords and Web-site information as files in encrypted folders. You need only remember one master password to open those folders. The Clipboard, normally used for the copy and paste, is completely erased when you close the program, which prevents its contents from being read or misused. For more details, consult the Appendix and Bonus Chapter 2.

1. Open Internet Explorer and choose Tools ⇨ Internet Options.

The Internet Options window appears.

2. Select the Content tab and click the AutoComplete button.

The AutoComplete Settings dialog box appears.

3. Uncheck the boxes beside the options you'd like to stop using, and then click OK.

Using the New Internet Explorer 7

Internet Explorer 6 (IE6) was released in 2001. By computer standards, that's pretty old. Internet Explorer 7 (IE7) is included as a staple with Microsoft Vista. It was released in late 2005 — not just in answer to Firefox and other competing browsers, but also to address growing security concerns.

You likely got IE7 in automatically through Windows Update, which method of distribution ensures that the security measures IE7 sets in place immediately affect the propagation of malware on the Internet via those machines — so long as it's downloaded and installed. Users have the option to download and install it; if you haven't yet, we suggest you do.



We highly recommend *all* Windows XP users upgrade to IE7, even if you use an alternative application as your primary browser. Keep in mind that IE is part of your Windows interface — Windows Explorer, for example, or any folder showing files and folders is using IE to display them regardless of what your default browser is.

What's new in IE7

IE7 has many new security enhancements, as well as a revamped user interface that's been modernized to include tabbed browsing — which puts it on par with Firefox and Opera. Microsoft made a highly acclaimed anti-spyware application available (Windows Defender) that's free to all. You can check out Windows Defender at the following URL:

```
www.microsoft.com/athome/security/spyware/software/default.aspx
```

What's new in IE7?

Getting started requires a slight learning curve — items have been moved to make them more visible, to accommodate tabbed browsing, and other new features. Once you've used IE7 for a little while, it's very intuitive. IE7 is a very capable and secure browser, equal to (or better than) any of the common alternatives for Windows XP and Vista computers.

Microsoft has implemented many new or updated security features in IE7 to help ensure your safety, and to help you with your routine security chores — these, for example:

- ✓ ActiveX controls are defaulted to “off” or disabled. ActiveX controls in previous versions of Internet Explorer were frequent targets for malware; the improved Add-on Manager gives you complete control over add-ons and allows you to delete unwanted ActiveX controls.
- ✓ The new active-scanning phishing filter monitors for known and potential fraudulent sites. It provides protection from phishing attacks, Internet fraud, and Web site spoofing, and will block sites if necessary. The phishing filter is optional, but we see no reason not to use it. It's reportedly updated several times an hour. The enhanced Security Status Bar changes color to get your attention. Phishing-filter notices and the secure connection padlock are now beside the address bar for easy reference.

- ✔ Cross-domain scripting is limited, thus minimizing unauthorized downloads, and other exploits.
- ✔ URL handling has also been redesigned to minimize exploits. If an unrestricted or trusted Web site attempts to import and execute a script from a restricted site, that script won't be allowed to run.
- ✔ IE7 will notify you if a Web site attempts to transmit your password information insecurely.
- ✔ IE7 offers a Protected Mode that enables you to surf with browser add-ons — such as ActiveX downloads, plug-ins, and browser-helper objects (BHOs) — disabled.
- ✔ The version of IE7 included in Windows Vista goes even farther by running in isolation from other applications — making it more difficult for your system to be compromised and turned into a threat. (The Vista version also integrates fully into Vista's Parental Controls.)

These are only a few of the newly incorporated features have brought Internet Explorer 7 to a level that's comparable to the other popular browsers, such as Opera and Firefox, which you can find at the following sites, respectively:

```
www.opera.com/products/desktop/  
www.mozilla.com/firefox/
```



One non-security-related feature added in IE7 is browser add-ons, which are programs intended to enhance the functionality of your browser. They may take the form of toolbars, pop-up blockers or just something to add more fun, but the bad guys can use them maliciously to control your browser.

IE7 is great but not a panacea

Just as dumping Internet Explorer in favor of Firefox could not guarantee you would be perfectly safe, neither can switching to IE7. You still need firewalls, antivirus scanners, spyware blockers, spam blockers, anti-rootkit scanners, and you still must keep them updated, be careful where you go on the Internet, and be careful of what you click. But that's no different than with any browser. It all boils down to practicing safe computing, which includes keeping our systems patched and upgraded with the latest security updates.

IE7 is not a cure-all for our Internet security woes, but it's an important step in keeping our computers, our personal information, and our children safe.



Using an alternative browser does not make you immune to malware or rootkits. Not by a long shot. As proof that no browser is immune to malware, the Gromozon rootkit actually executes differently depending on what browser is being used so it may adapt its MO on the fly; that is, it executes differently depending on if you have IE, Opera, or Firefox. Very sneaky!

Surfing with Firefox instead

We can certainly see why so many people love Mozilla's Firefox Web browser. It has many more features than are provided in IE7. Mind you, Firefox came on the scene long before IE7, and Firefox 2 was released just days after IE7. It's an excellent browser that's capable and quick, with security enhancements that are easy to set up using one of Mozilla's free tutorials. Focusing on usability and features, Firefox has put a large dent in the desktop Internet browser market, enough to make Microsoft sit up and take notice. Some of the main improvements and features of Firefox 2:

- ✔ Setting up Firefox 2 is a snap. Select Tools ⇨ Options. The Options window appears. Across the top are icons for each aspect of the browser. Click them to view and change your settings. Very sleek.
- ✔ The Help file is quick and easy, including a section for Internet Explorer users that helps them understand the nomenclature used in Firefox compared to how it looks in IE. Nice touch.
- ✔ Enhanced security has been added to block phishing scams and warn you of possibly fraudulent Web sites.
- ✔ A built-in search box (on the upper right) with default engines you can use, or you can add on your own particular favorites. You can highlight a word or text in a page and drag it to the search box, or you can type in an entry if you wish. When you type into the box, a drop-down list will appear with options for your query. Select one to view the search page corresponding to it at the search engine of your choosing. You can get more search engines at the Mozilla Web site's Add-ons section (see the last bullet in this list).
- ✔ An RSS News Reader that you can use by clicking the Latest Headlines button in the upper left. To add feeds to the Bookmarks Toolbar Folder click the icon that indicates a feed is available in the address bar (same icon as beside the Latest Headlines). The new feed is then automatically added to the list. Firefox calls these feeds Live Bookmarks.
- ✔ Tabbed browsing allows you to open multiple pages in a single window, and you can re-order their positions with drag-and-drop. Now you can have more tabs than will fit in the viewing window with scroll arrows that appear on either side. If you accidentally close a tab, you can retrieve it from the "Recently Closed Tabs" list in the History menu. Firefox even forces Web sites that try to open in a new window into a tab instead, saving time and resources on your computer.
- ✔ Getting add-ons for your browser is now simple and direct; just choose Tools ⇨ Add-ons. A box will pop up showing two icons for Extensions and Themes. Click Extensions, click Enable, and then click Get Extensions at the bottom of the window. This opens the Firefox Add-ons

page in your browser. From this page you can select and download from hundreds of add on tools and plug-ins. Once added to your browser, Firefox will regularly check for updates to them.

Comprehensive lists of Firefox add-ons, plug-ins, extensions, themes, and search engines are freely available at the Mozilla Web site in the Add-ons section. These allow you customize your browser the way you want it. You can check out these cool features at the following URL:

```
http://addons.mozilla.org/firefox/
```

You can download Firefox 2 here:

```
www.mozilla.com/en-US/firefox/
```

Installation instructions are featured at this page:

```
www.opensourcearticles.com/introduction\_to\_firefox
```

By all means take Firefox for a test drive and see for yourself why so many people are now using it as their primary browser. No application is without fault, but Firefox has proven its worth. Patches and updates come out rapidly, when they're needed most. We like it!

Staying ahead of the game with SiteAdvisor

SiteAdvisor (SA) is a program that can help you evaluate a Web site's safety before you even you dare to tread there. The program was developed by MIT Security Researchers and was recently acquired by McAfee, Inc.

SiteAdvisor works with both Internet Explorer and Firefox. When you do a Google search, if you hover your cursor over the SiteAdvisor icon tagged on to your search hits, you can view SA's recommendations about how secure (or not) a site is without even having to venture there.

Sites are rated using a simple color-coded system to indicate SiteAdvisor's recommendation action:

- ✓ **Red:** Avoid. Serious issues to be considered.
- ✓ **Yellow:** Proceed with caution. Reservations are indicated.
- ✓ **Green:** Go for it. The site was tested to be safe.

You can download SiteAdvisor here:

```
www.siteadvisor.com/download/ff_learnmore.html
```

Questions considered by the Site Advisor developers in establishing a rating for a given Web site include the following:

- ✓ Do site downloads contain malware or bundled undesirable programs (such as adware and spyware)?
- ✓ Are you exposed to a pop-up barrage or nuisance-type behavior at the Web site?
- ✓ How many spam e-mails are generated if you complete and submit a sign up form on the Web site?
- ✓ What is the customer/client feedback from SiteAdvisor users who have visited the site?



SA observes the principle of informed consent. You, and only you, decide what to do, basing your decision on the information it provides. That way, if a site is listed as red and you want to visit it anyway, you have the option.



If you don't want to install SiteAdvisor, you can enter a domain name to retrieve information on a specific Web site without installing the tool here:

```
www.siteadvisor.com/sites/galttech.com
```

Must-Have Protections Online

To best prevent malware getting your system into its vile clutches, you need to protect it with software made for that purpose. Specifically, you need to

- ✓ **Install a firewall:** A firewall is your first line of defense against malware and hacking attempts. Your computer has 65,535 *ports* — points where contacts can be made — more ports than any coastline on the planet. To protect them, you need a good quality firewall that's regularly maintained and updated as needed.
- ✓ **Install Anti-Malware Scanners:** Although most scanners cannot detect nor remove an active rootkit, they can stop the malware that installs them. A criminal hacker needs to use a trojan or a backdoor to get a rootkit into a computer. Using security programs that offer excellent onboard active protection will greatly improve your resistance to rootkits.

The following sections get you up to speed on what firewalls and scanners are, reviews the various types of each, and lets you know which you need to protect your system.

Firewall first

Firewalls are an absolute must online. Without one, your system is wide open to most trojans and simple hacking attempts. Hackers regularly run search programs online, which they use to probe for open ports and software vulnerabilities. It's so easy to compromise an unfortified system, which can then be used to hack into something bigger and juicier without being traceable back to the hacker. The owners of those PCs could be left "holding the bag," held legally responsible for crimes they didn't know their compromised equipment were committing. These compromised systems become a menace to the rest of us once they have been transformed to worker bees for the enemy. Below we explain what a firewall is, and what it does, featuring software and hardware versions. Routers are included to show you their relationship to hardware firewalls.

What is a firewall? It's either hardware or software that surrounds your system with a shield, limiting access to potentially dangerous, unwanted or hidden communications between your system and the Internet. Well-designed firewalls have rules that are both automatically set and manual settings that the user can customize to their needs.

Understanding how firewalls work

To understand how a firewall works, you first need to understand how your system interacts with the Internet. Let's give a simple example: You want to get your e-mail. So you open your e-mail client such as Outlook Express, Outlook, Thunderbird, and so on, and click the button that downloads your e-mail from your e-mail account. What's actually happening when you do this is that your e-mail program is sending a request out to your e-mail server — the online location that actually receives your e-mail and holds it until you're ready to retrieve it. That request says to the e-mail server "Hey, it's me, and I want to get my e-mail." The e-mail server gets that request, makes sure that the username and password are correct, and then sends back all your e-mail to your e-mail software.

Behind the scenes, your request is sent as a "packet" of information that your e-mail server needs in order to know that you want your e-mail, as well as your identification, and your IP, which is your current address on the Internet. With the information in that packet, your e-mail server knows that you are who you claim to be, as well as where you are on the Internet — and then sends you your e-mail.

There are 65,536 ports, ranging from 0 to 65,535, but only a small number of them are actively used at this time. There is an international organization that's responsible for assigning ports for specific uses, the Internet Assigned Numbers Authority (IANA). Most home users only use a very small number of assigned ports. The two ports used for sending (Port 25) and receiving (Port 110) e-mail are used by almost everybody. In addition, port 80 is the port used for requesting and receiving Web pages.

A normal home user may only use less than 25 ports out of the 65,536 possible ports for normal daily communications on the Internet. The rest of the ports aren't generally needed, so it's very safe and desirable to block those ports from communicating with the Internet. In fact, if you don't block those ports, they may be used by malware without you even knowing it's happening. Blocking unused ports from outbound and inbound communications with the Internet is one of the most important functions performed by a firewall. You also want to restrict what software can use the ports you do need. Blocking or permitting software from using needed ports, or from the Internet entirely, is another important firewall feature. We will discuss this more below and introduce additional firewall functions as we move along.

Two major types of firewalls exist: hardware firewalls and software firewalls. *Software firewalls* are installed just like any other program on your system. *Hardware firewalls* are embedded into hardware. The following sections describe each type in more detail.



Our strong recommendation, considering the major role firewalls play in your security suite, is that you always use a software firewall, at least for its program control features, and if you use broadband, use a hardware one as well. Doing so greatly improves your computer's ability to stay secure.

Hardware firewalls

A *hardware firewall* is a small, low-power computer in a separate box that sits between your broadband modem and your computer or computers on a LAN. The correct name for this small device is a *hardware router*, and to understand how it works we need to examine both parts of this specialized small computer to see what they do to protect your big computer from malware threats, exploits, and hackers.



Anyone using dial-up can ignore this and skip to the “Software firewalls” section a bit later in this chapter. Generally very few hardware firewalls work with dial-up connections; there are a few that do, but they're expensive and hard to find. Due to the nature of how Internet addresses are allocated by your Internet Service Provider (ISP), hardware firewalls are less often needed. If you use dial-up to connect to the Internet, you get a new Internet address each time you connect. That makes you harder to find and less likely to experience inbound threats from the Internet, as long as you have other protections in place as part of your security suite.

Unlike dial-up connections, broadband connections are assigned more or less permanent Internet addresses by their ISPs. And broadband is always on. That makes it easier for you to be found — and makes a hardware firewall even more important.

We recommend that you consider getting yourself a router or hardware firewall to increase your security online. Routers that include a hardware firewall as well as standalone hardware firewalls are commercially available at reasonable prices. Dedicated hardware firewalls are commercially available. Here are a few links to hardware routers and firewall sites:

```
www.dlink.com/products/category.asp?cid=2  
www.cisco.com/en/US/products/hw/routers/index.html  
www.sonitwall.com/products/index.html  
www.netgear.com/products.php
```

Why use a hardware firewall? Very simply, if a hacker breaks through the hardware one, all he gets for his trouble is empty space. The hardware has no programs he can use to access your computer software. Even if hackers manage to breach the hardware firewall, nothing exists just past it to modify or grab onto, and a little further along they meet up with the software firewall.



The not-so-good news is that though hardware firewalls definitely thwart computer break-in attempts, they provide no protection from rootkits or any other type of malware acquired by visiting a malicious Web site or from installing a piece of software. Neither can they protect you from malware that secretly launches itself from within a legitimate system process. This is where the active protection components of your anti-malware programs should kick in.

Adding a hardware router or firewall to your Internet connection is a straightforward process. Since your system already connects to the Internet, all we need to do is change the wiring a little bit. Your computer already is connected to a broadband modem. The modem is connected to either your cable television line, or to your phone line, depending on whether your ISP provides Internet connections via cable or by phone line (called DSL or ADSL, depending on the ISP). The only change you need to make is to disconnect your computer from your modem, connect it to your new hardware router or firewall, and then connect the hardware router or firewall to your modem. That's it! In many cases the standard internal settings the hardware manufacturers set as defaults will let you connect to the Internet immediately without any further actions on your part. You should read the "Quick Start" instructions that come in the box with your hardware router or firewall, as that will have both pictures and instructions on exactly what sockets to use to connect to your modem and your computer.



Several manufacturers are now starting to produce compact portable hardware router or firewalls for frequent travelers who often find themselves connecting to the Internet from non-secure locations. Within the last year or two, several very tiny units, and thus easy to carry and toss into your travel gear,

have been brought to market. These units do possess the usual range of port stealthing, SPI and wireless protection of full-sized hardware routers or firewalls in a very compact package. Again, these units are not expensive, and are readily available for less than \$100. But these units are not for dial-up connections.

All modern home-grade hardware routers or firewalls have browser-based interfaces. What this means is all you need to do to access the settings in your hardware router or firewall is to enter a specific IP into your browser and tap enter. The “Quick Start” instructions will tell you what IP to use, although many use a standard one, such as 192.168.0.1, for example. Once you access the correct IP for the hardware router or firewall, you should see a login screen requesting a name and a password. Again, the “Quick Start” instructions are your friend, and will tell you the default password for your hardware. Often times it uses a simple username like “admin” and the password is “password,” or blank by default. Be sure to change the username and password after saving your settings. Hackers know that many units use standard usernames and passwords, and will try to exploit that standard setting.

Another very important thing to check in the “Quick Start” instructions (or in the accompanying manual) is how to reset your hardware router or firewall to its default settings. That information will save you when you have forgotten your password, or entered some setting that locks you out of the unit. Usually, your unit will have a reset button on its back panel. But if you reset the unit, doesn't that also mean that all your custom settings are gone? Not if you have done it the right way and exported your settings to a file located on your hard drive. All modern units have a Settings Export and Settings Import feature, and you should always keep the current settings file stored on your hard drive.

Now let's examine some of the features of your new hardware router or firewall, starting with the router.

Simply put, a router is a hardware network device that permits a number of individual computers to look like a single computer to your modem and ISP. Using a router is the easiest way for several computers in a home network to easily share a single connection to the Internet. And it's very efficient in doing so. You probably won't even notice any speed difference whether one, two or even more individual computers are working on-line at the same time. Sharing a single Internet connection using a router is easier than trying to set up Windows Internet Connection Sharing.

So what makes a hardware router into a firewall? Several additional true firewall components, including network address translation (or NAT), port blocking, port stealthing, and stateful packet inspection (or SPI). Let's examine these in turn.

Network address translation (NAT)

How does a router manage to do this? It uses a technology called Network Address Translation (NAT). NAT permits the router to use a Dynamic Host Configuration Protocol (DHCP) to assign individual and unique internal IPs to each computer on the router's LAN side, while taking over the IP assigned by your ISP on the router's WAN side. It then keeps track of which computers are working on the Internet from its LAN side and then knows how to route return packets back to the correct computer. The router does this by using a routing protocol and routing table to keep track of, and to convert external and internal IPs in packets that flow through the router. In simple words it acts like an old-fashioned "party line" telephone where multiple users can use a single telephone line for shared communications, only it's a little more sophisticated than that — it keeps each computer's communications private from the others on the same LAN.

One effect of NAT is that your router acts like a firewall to an extent. This is because the external IP is the only one seen by the outside world, and the outside world cannot either see or send packets directly to any computer connected on the LAN side through the router, unless the router knows where to route those packets on the LAN, which it cannot do if those packets were not requested by a computer resident on the LAN side of the router. While NAT was neither designed as, nor intended to be a firewall component, it does provide significant protection and prevents the external Internet from reaching computers on the LAN side of the router.

Don't be fooled, though. It's possible to bypass NAT, but it takes a very clever cracker and some time to accomplish that feat. Don't assume that a router will provide full firewall protection to your computers; it won't. Many manufacturers claim firewall features for NAT, but that's advertising, not fact. NAT is not a true firewall, nor will it ever be by design. It's an inherent requirement of routing; nonetheless, it does help protect your computers and also add valuable protective features.

Port blocking and port stealthing

Depending on the capabilities of your unit, you can selectively decide which of the 65,536 possible ports you want to block on the unit's WAN or Internet side. Generally, you should block all ports unless you absolutely need them for your online services. Most hardware router or firewalls come pre-configured to do just that, and it usually isn't necessary to alter those default configurations unless you find that something you need just doesn't work. Many better home units also permit you to block all unneeded outgoing ports as well, which is a very good idea. Again, those settings are usually the unit's default settings. Even better than port blocking is port stealthing.

When someone wants to see if a computer is using a particular IP, they send the target computer a packet called a *ping* — which is similar to a sonar ping. With sonar, you send out an audible ping, and the sound wave is echoed back to your sonar unit by the target, allowing you to see whatever is there. So too in computer communications: If you want to see if “anyone is home,” you send out a ping, and if it’s answered, you know there is a computer on the other side.

If you want to try it yourself, it’s easy to do; follow these steps to ping `www.google.com`:

1. Click Start and choose Run.

The Run window appears.

2. Type `cmd` into the **Open: field and click OK or press Enter.**

The command prompt appears.

3. Enter `ping www.google.com` (making sure there is a space after the word ping) and press Enter.

Your results should look like this if you’re connected to the Internet:

```
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985–2001 Microsoft Corp.

C:\Documents and Settings\{your_username}>ping www.google.com
Pinging www.l.google.com [216.239.37.99] with 32 bytes of data:

Reply from 216.239.37.99: bytes=32 time=15ms TTL=240
Reply from 216.239.37.99: bytes=32 time=16ms TTL=240
Reply from 216.239.37.99: bytes=32 time=16ms TTL=240
Reply from 216.239.37.99: bytes=32 time=17ms TTL=240

Ping statistics for 216.239.37.99:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 15ms, Maximum = 17ms, Average = 16ms
```

Well, Google’s there all right — would you expect otherwise? And its gateway computer told us in no uncertain terms that it’s there. If it wasn’t there then we would have seen four messages saying that the ping had timed out — in other words, either no one was home, or else someone really clever was home but refused to acknowledge the ping.

That’s what port stealthing does — it creates a “black hole” where pings come in, but never come out. With only port blocking, but not stealthing, each ping would have been answered, and whoever pinged you would know you were there, even though the port was otherwise blocked.

Most modern home-grade hardware routers or firewalls do port stealthing, and that feature is usually set by default. But if you see a control that talks about returning pings from the Internet, port stealthing, or something like that, check your manual carefully to make sure that it's properly set to stealth your unit on the WAN side.

Fortunately, there is a great and easy way to check whether you have successfully blocked or better yet stealthed your ports. Here's how you do that.



Enter **www.grc.com** into your browser. Wait until the redirection happens and takes you to the home page of Steve Gibson's excellent and very informative Web site. Look around a bit; there is some very useful information there on the Internet in general and more specifically about how to secure your systems properly. Page about midway down and you see a link to ShieldsUP! Click that link and it will take you to new page. Toward the bottom, you see your true external IP (now, don't panic about that — it's a natural part of Internet communication) and a couple of boxes marked "Proceed". Click one of them; on the page that links to, choose the All Service Ports test. (You can also try the other tests he offers; they're all very helpful.) The test will start; it checks all ports from 0 to 1055, which contains almost every common service port a home user would ever need. What you want to see is a sea of green. If you do, congratulations, you're totally stealthed. If you see any blue, those ports are not stealthed, but at least they're blocked. Any red indicates the presence of an open and unprotected port — and that's a vulnerability that you need to close.

Stateful Packet Inspection

This is the second major firewall component, and it's a critical one. If a unit doesn't say it supports Stateful Packet Inspection (SPI), don't buy it. SPI is like a "one-way swinging door." If your firewall has SPI, it remembers the state or nature of requests that have been made from computers on the LAN side of your hardware router or firewall. Suppose you request your e-mail; SPI will permit entry only for those packets it stored the state for — or rather, for those packets received from your e-mail server in response to the request for e-mail. Any unrequested packets (perhaps spoofing a response to a fake request for e-mail for example) won't be permitted entry because there was no corresponding outbound request. Thus SPI significantly increases protection for allowed ports beyond even the protections afforded by Port Stealthing. Nothing will be permitted inbound unless there was a specific outbound request preceding the inbound packets. Most better-quality home-grade hardware router or firewalls support SPI for the home user.

How much does this magical hardware cost? Not much at all. At the time this was written, many good-quality home-grade units supporting all the recommended features could be found either on sale in local computer shops, or

online, for between \$50 and \$100. When you consider just how much protection a good hardware router or firewall adds to your Internet security, it's almost impossible not to get one.

Many broadband ISPs these days supply combination units to their customers that are a combination of a broadband modem and a hardware router or firewall, all built into the same unit. It makes economic sense for the ISPs to do this, since the cost to their customers of going unprotected — and all the service calls that will lead to — far exceeds the additional cost of using these combination units over the cost of a plain unprotected modem. Furthermore, many customers don't understand how to set up the components of a unit like that on their own. Yet these customers still have the right to access the Internet, so the ISP has an obligation to try to protect them somehow. If you're already on broadband, you should check with your ISP to see if your modem already includes a router and firewall. If not, many ISPs will replace older modems with combination units for a small fee — or free, as part of their service contract.

We have spent some time discussing what a hardware router or firewall can do for you. Now we need to tell you one major thing it won't do: It *won't* block outbound Internet access, depending on the software requesting the access. Consider your hardware router or firewall as a separate appliance on your network. Suppose it sees an outbound packet using port 80, these packets request Web pages, and are usually created by a browser or e-mail client. If malware attempts to access the Internet using a permitted port (80, for example), there is no way for a hardware router or firewall to tell the difference between a legitimate outbound request (from your browser, say) and an outbound request on port 80 by malware. If that outbound request is permitted, SPI remembers the state of that packet; in response, it permits inbound packets to access the LAN and be sent to the computer making that request. The only way to prevent malware-generated outbound requests from accessing the Internet is to use a software firewall that runs on every system on the LAN.



Software firewalls

For workstations and standalone computers, examples of excellent firewalls are available on our CD. See the Appendix and Bonus Chapter 1 for more about how to use them.



For network servers, the security applications and hardware available from the following sites are excellent, economical, and recommended by CastleCops experts:

```
www.watchguard.com/products/zeroday.asp  
www.watchguard.com/products/utm.asp  
www.sonicwall.com/totalsecure/index.html  
www.winternals.com/Products/ProtectionManager/
```

Two kinds of software firewalls exist: rule-based and application-based.

- ✓ *Rule-based firewalls* have gotten a little easier as guides have been developed by which you can configure them. In the old days, you had to write your own rules — a daunting task for new users.
- ✓ *Application-based firewalls* have their rules already set up. All you have to do is decide what applications you want to give Internet access and how high to make the security settings. Piece of cake.



For effective security, you can use only one software firewall at a time. Two software firewalls running in tandem may “collide” when they both attempt to monitor the same events — and the results can be very unpredictable. We also mention this here — not because we’re being pedantic but because it’s such a critical point, especially given that Windows XP SP2 enables the Windows Firewall by default. Users of the ZoneAlarm firewall will find that enabling ZoneAlarm automatically turns off the Windows Firewall, saving them the extra step of disabling that feature. Note that Agnitum Outpost and Sunbelt Kerio Personal Firewalls do likewise.

Following are some links to software firewall providers:

```
www.zonelabs.com/store/content/home.jsp  
www.agnitum.com/products/outpost  
www.sunbelt-software.com/Kerio.cfm
```

How software firewalls work

A *software firewall* is a set of programs that provide port-blocking features similar to those of a hardware router or physical firewall. Software firewalls don’t provide any routing controls — nor (for the most part) do they provide SPI protections. (Although a couple of software firewalls do include SPI, they’re expensive and intended primarily for professional and commercial use.) The major difference between hardware and software firewalls is that software firewalls can control both inbound and outbound access, based on which software receives or creates the packet(s). If you connect to the Internet using broadband, that aspect of a software firewall is an important complement to the protections afforded by a hardware router or physical firewall. If you connect by dial-up, a software firewall may be the only practical way to add firewall protection to your security suite.



If you often use the laptop to connect to the Internet from hotels and such, you must use a software firewall to protect yourself when you’re not behind a hardware router or firewall set to your individual security needs.

Software firewalls work by using several components deeply embedded in the operating system’s kernel (similar to a rootkit but beneficial). These intercept inbound and outbound packets before they’re processed — and before they’re permitted into or out of the higher levels of your system and software. In fact,

many rootkit-diagnostic tools report the kernel hooks that software firewalls add to your operating system as if they were kernel-level rootkits. Note, too, that some other anti-malware and antivirus programs add kernel hooks. Because software firewalls embed themselves so deeply into the operating system, they can monitor every packet and “see” what software is trying to send the packet or is “listening” for it to be received — and can block these packets if a non-permitted or unknown program is sending or listening. This process is known by various names but can all be lumped under the general term of *software rule-based packet control*. It’s rule-based because for individual software, specific rules can be created within the software firewall to permit or block specific applications for Internet access by port and also by IP if necessary. Software firewalls will block by default any application for which a rule is not created. Most software firewalls also support port stealthing by default.

Setting firewall rules

Firewalls use settings called *firewall rules* to know what and what not to block. If your firewall warns you about an application and you tell it to allow that application to connect, the firewall sets a rule for itself so that it knows to allow that application to connect from henceforth. One might think that rule making is a very difficult task, but in fact, it usually isn’t. Most good software firewalls have preconfigured rules for major browsers, e-mail clients and also for most components used by your operating system that require Internet access to work properly. If a program for which there is no preconfigured rule attempts to access the Internet, good software firewalls pop-up a warning, telling you what software is attempting access, and asking you to decide whether to allow the access or not. Generally, you can permit a single time access, or set a rule for permitting future regular access, or do the opposite deny access — either on a one-time or regular basis. Those rules are saved by the software firewall and will be applied automatically each time the software or application tries to access the Internet.

Since every computer is fairly unique in terms of the software it uses, every software firewall will take some training to configure appropriate rules for each software component that requires Internet access. Usually this is a somewhat annoying process, requiring dealing with frequent warning pop-ups from the firewall, but isn’t hard to do, nor does it take much time. You can accelerate this training by simply starting all your most used software, respond to the firewall pop-up to create an appropriate rule, and then close the software subsequent to creating the rule. Most rules are simple, the firewall sees the software trying to access the Internet, asks if it should permit or deny it once, or always. A simple response to the pop-up automatically creates a new rule for your firewall to use in the future. Once a rule is set, you should never see another warning pop-up unless that particular software component is changed for some reason.

One safety mechanism built into most firewalls is that it looks for changed components as well as new ones, and when a component changes, it triggers a pop-up to let you know that the component has changed. If you have just completed an update for that software component, well, it's obvious why the firewall pop-up has been triggered. But if a pop-up is unexpectedly triggered, and you have not done a software update, that may be a sign that some malware is spoofing a known common software component.

Another point to be aware of is that software firewalls are frequently updated. These updates occur for many reasons, including fixing bugs in the software, additional security components or features, conflicts with other software, and so on. So you should check for updates to your software firewall fairly regularly, and be sure to keep it up to date. Most current software firewalls will provide you a means to obtain automatic updates.

The Windows Firewall

Windows XP comes with a built-in software firewall that is turned on by default. The Windows Firewall is adequate for the basic security needs of a simple user — but if you want something more powerful and effective, get a commercially available firewall.

Why the Windows Firewall leaves something to be desired

With so many bidirectional firewalls commercially available, there are certainly more effective alternatives to the unidirectional Windows Firewall. The primary reason many consider the Windows Firewall inadequate is that it won't stop outbound connections to the Internet. Some viruses and trojans can disable it completely, but other firewalls can also fall prey to that malware trick. What's important about the Windows Firewall is that it provides basic onboard protection for Windows users who would not ordinarily pay much attention to installing a firewall. The new Windows firewall, Automatic Updates, and the Windows Security Center were SP2 improvements aimed at making security solutions available to all Windows XP users (especially those who otherwise might not have taken the initiative to seek a security solution themselves). Many popular commercial firewalls configured to work with Windows XP SP2 surpass the Windows Firewall in functionality; some very effective ones are free to home users. The Windows Security Center can monitor most of these third-party firewalls, just as it can monitor antivirus functionality. Another option is to subscribe to Windows Live OneCare, which does include a bidirectional firewall as part of its complete security package:

```
www.windowsonecare.com/prodinfo/Default.aspx/?sc_cid=sah
```

Because the Windows Firewall can only prevent *inbound* intrusion attempts, we strongly recommend using a *bidirectional* software firewall to restrict both inbound and outbound traffic. A good firewall provides protection on both

sides of the door, which prevents intrusive programs from phoning home and transmitting personal information about you, your computer, or your browsing habits to remote servers (a practice most users would frown upon). Anything wanting to leave or enter has to check in with the firewall first.

As an example of why you need outbound protection as well as inbound, say you download an application without checking to see whether the source is trustworthy — and (oops) it comes with a trojan bundled with the installer. After you install the application on your machine, the trojan attempts to communicate online with its master, the criminal who created it. The Windows XP Firewall won't notice this activity. The trojan records and sends your passwords, financial data, and access codes to the criminal — while you proceed with your business, blissfully unaware that your PC is being hacked. Of course, this is a worst-case scenario; it assumes that your security program's active protection was unable to curb the threat. But that's not entirely out of the realm of possibility — and it's actually a distinct possibility if the rest of your protection isn't up to snuff.

Turning off the Windows Firewall

If you install another software firewall, do be sure to turn off the Windows Firewall before you install and use an alternative program.



Before you turn off the Windows Firewall, be sure to have an alternative product. The Windows Firewall is preferable to having no firewall at all.

Here are the steps for turning off the Windows Firewall, but first . . .



Do the following *while you're offline*, if possible. Going offline to make changes is important; if you stay online, you have no protection for a brief time. Yes, you're only turning off one firewall and then enabling another one — but malware can invade your computer in less than a millisecond.

- 1. Click the Start button and choose Control Panel.**

The Control Panel appears.

- 2. Double-click Security Center (if your Control Panel is in Category View, it's a single click).**

The Security Center window appears.

- 3. Click Windows Firewall at the bottom (under Manage Security Settings for:).**

The Windows Firewall window appears.

- 4. Click the General tab, where you see green and red security shields; click the red shield with the "X" in it (to turn off the Windows Firewall) and then click OK.**

Now install your alternative firewall and reboot. If the software firewall you've chosen is configured for Windows XP, the Security Center will then show that the firewall is On.

Knowing what firewalls you need

When deciding on what firewalls — software and/or hardware — you need for your system and/or network, keep the following points in mind:

- ✔ If you're a home dial-up user, a software firewall is all you can use; dial-up doesn't work with routers. If you have a software firewall but want extra security — and you use cable or high-speed DSL (or wireless connections) — then get a router or hardware firewall, or both.
- ✔ Even if you don't use some form of hardware firewall, you simply *must* have a software firewall if you're doing anything much online. Many personal firewalls are free of charge for noncommercial use, so you have no excuse not to have one.
- ✔ One good-quality software firewall is all you need. More than one will cause system conflicts — and be worse than none. The only case in which two firewalls work is if you use a hardware firewall alongside a software firewall.
- ✔ We discuss various quality firewalls, including ZoneAlarm — widely considered the world's best software firewall — in Bonus Chapter 1. ZoneAlarm notifies you whenever anything wants access to your computer — and you can block the intruder completely, either by having the program block access automatically or by manual selection. ZoneAlarm Professional provides state-of-the-art protection from rootkits.
- ✔ A router is a must for wireless networks, especially when it comes to protecting small business networks in general.
- ✔ If you use a wireless network, you absolutely need a software firewall — in addition to the hardware firewall included with your wireless router.



For a more detailed explanation of how firewalls work, and which one is right for you, have a look at *Firewalls For Dummies* by Brian Komar, Ronald Beekelaar, and Joern Wettern (Wiley Publishing, Inc.).

Scanners Next

After you've decided upon — and installed — your firewall solutions, the next most essential security measures are installing antivirus, anti-trojan, and anti-spyware programs on your system. Then, when a virus, trojan, or spyware interloper is detected, the appropriate program will move the infected file to a quarantine area for disinfection or removal, preventing the malware file from making contact with any other program. As long as you keep the virus,

trojan, and spyware data files up to date (check every day), detection is reliable. Apply all updates and program patches *as they're released*. Most applications do these update tasks automatically, or you can schedule them yourself.



No antivirus, anti-trojan, or anti-spyware application is ever 100 percent effective. There may be times when your application won't remove, contain, or stop a particular threat.



Some of the best antivirus, anti-trojan, and anti-spyware scanners are available on the CD. See the Appendix and Bonus Chapter 1 for more information.

Three types of scanners . . . sort of

Different types of scanners have been developed to address each type of malware, but some of them overlap. Most scanners are dedicated to a specific type of threat. For example, antivirus (AV) scanners address viruses, anti-spyware scanners address spyware and adware, and dedicated trojan scanners exist to address trojans (as Trojan Horse programs are known these days). Most antivirus programs will also detect viruses, worms, and many trojans.

As the borderlines between malware classifications become increasingly blurred, security programs have expanded their threat coverage accordingly. Here are two examples:

- ✓ ewido, formerly billed as an anti-trojan program, is now known as AVG Antispyware.
- ✓ Panda Active-Scan detects both spyware and viruses, though it removes only the viruses.

Another reason scanners are incorporating extended coverage is to deal with the many new threats of a *blended* nature — malware with multiple components that span more than one category.



Even though the distinction between what the different types of malware scanners will detect and remove has become increasingly blurred, it's still advisable to employ a variety of different scanning programs to obtain maximum protection. Individual scanners are still best at doing what their name implies or whatever category of scanner they fit into.



In addition to the real-time scanning done by your applications, run their full system-scan capabilities at least once per week — preferably offline, as many malware programs need the Internet to maintain themselves. Turn off all nonessential programs and run full scans with your antivirus, anti-trojan, and anti-spyware scanners, one at a time. This may take some time, depending on the size of your hard drive(s) and the number of files you have.

To stay healthy, your system needs full scans on a weekly basis, whether you're on a network or not. You may scan more often, depending on how much surfing you do. We recommend full scans, using all your regularly updated scanning tools, *at least* once per week while your system is offline.

The following sections discuss each of the three types of malware scanners in detail.

Anti-spyware software

Anti-spyware software is fast becoming another must-have for online activities, protecting you from the scourge of spyware and adware that can otherwise rob you of your privacy, identity, and well-being. Spyware can track you online, record your keystrokes (*keyloggers*), and steal your passwords and PIN numbers. Adware can track your surfing habits and force you to view pop-up ads based on them, or hijack your browser search function, forcing you to view the Web sites that they choose. Anti-spyware tools are included on the *Rootkits For Dummies* DART CD.

You need at least one program that scans for spyware and gets rid of it — and it should include an active-protection component so it can be updated automatically, and complement the real-time protection that your AV provides. Many freeware anti-spyware programs are excellent on-demand scanners but don't include real-time active protection. Still, don't let that dissuade you from using them. Using more than one on-demand anti-spyware program will expand your threat coverage.



Anti-spyware applications will scan for spyware, malicious adware, and hijackers. Real-time monitors scan the active files on your computer for infections, both on and offline, but they don't provide updates to the signature files that identify new malware.

Here are some links to popular and reputable anti-spyware vendors:

```
www.webroot.com/  
www.sunbelt-software.com/CounterSpy.cfm  
www.pctools.com/spyware-doctor/
```



Be especially cautious when you purchase commercial anti-spyware applications. A great many bogus and rogue programs out on the Internet exist solely to bilk the public out of their money. They know you want to protect your computer from spyware — so they're eager and willing to "help" you with applications that either don't work or are a form of spyware themselves. To view a list of rogue or suspect anti-spyware applications, browse to this security Web site:

```
www.spywarewarrior.com/rogue_anti-spyware.htm
```

Antivirus software

Antivirus software is absolutely required for any and all online and offline activities. *Antivirus programs* are designed to detect and dispose of viruses, providing both active protection and on-demand scanning (scanning individual files/folders or your entire computer for viruses). Your antivirus program should also include the capability to update the scanner with a real-time monitor — which serves as your active antivirus protection while you're online and offline (should you unknowingly open a malware threat contained within a zipped archive).

Antivirus applications are becoming more anti-trojan as well, covering many of the more popular versions of these malware. Trojans are not viruses, but they're proliferating rapidly. The vendors of antivirus applications are sensitive to the needs of their customers and so provide accordingly. Get an antivirus application, install it (after first making sure your system meets its requirements), and keep it up to date. Regular — daily — updates are essential to maintaining the strength of your antivirus application.



It's important to have only one antivirus active-protection component “on” at any time, otherwise conflicts and computer problems may arise. However, there is no problem in using more than one antivirus or anti-spyware on-demand scanner.

Antivirus programs provide real-time monitors (RTM) for while you're online. These scan for the full list of the most popular viruses.

Following are some links to some reputable antivirus software vendors:

```
www.grisoft.com/doc/1  
www.eset.com  
www.avira.com  
www.kaspersky.com
```

Anti-trojan

A program that scans specifically for trojans — preferably one you can update automatically. Trojans are not viruses; scanning for trojans requires different parameters. Usually anti-trojan applications scan for backdoors and keyloggers as well. As with antivirus scanners, purchased anti-trojan applications include a real-time monitor; freeware versions don't.



The essential difference between viruses and trojans is viruses propagate themselves while trojans must be downloaded by clicking a link or by opening an e-mail or an e-mail attachment.



Some antivirus vendors even supply tools to combat specific Trojan Horse threats. For example, Symantec has a well-known online library of malware-removal tools. It offers programs created to remove specific threats that a general antivirus scanner alone can't eradicate. You can find this excellent resource at

www.symantec.com/avcenter/tools.list.html

Even though antivirus software successfully disinfects most trojans, dedicated anti-trojan scanners are available — and they make a highly effective addition to your protection. Employing both antivirus and anti-trojan programs means greater coverage of more threats, and increased resistance to the malware that installs rootkits.

Here are three links to reputable vendors of anti-trojan software:

www.agnitum.com/products/tauscan
www.moosoft.com
www.misec.net

Knowing which functions you need in a scanner

Among the most useful scanner features for fighting rootkits are these:



- ✓ **Auto-updating feature:** This mechanism enables you to have new malware definitions installed automatically as soon as they become available.
A scanner is only effective if its definitions are updated to protect you against current threats.
- ✓ **Autoscheduler:** This feature allows you to specify when — and how frequently — you'd like a scan performed.
Both the auto-update feature and autoscheduler lend themselves to the “set it and forget it” philosophy — which is not only a desirable convenience to have, but also an ensured form of protection.
- ✓ **On-demand:** *On-demand* means you determine when to perform a scan, and you run the scanner when it suits you. Most antivirus and anti-trojan scanners enable you to scan selectively — whether it's a file, folder, entire hard drive, or other storage device. Many on-demand scanners enable you to scan a single file from a context menu by right-clicking the file you're interested in scanning.
- ✓ **Active protection (real-time protection):** This component scans for file-read and file-write operations in real time, with the goal of catching invaders in the act — which immediately prevents threats from installing malware or causing an active infection.



✔ **Heuristics features:** *Heuristics* tries to detect new malware variants by using generalized signature scans setup to detect “families” of viruses or trojans. If a file being scanned appears to be like a particular virus or trojan family, heuristics will detect it as a suspicious or unknown threat. Enabling the *heuristics* feature in antivirus and anti-trojan scanners is important but be aware that it can cause more false positives. It detects possible viruses or spyware, depending on the scanner. Set it to scan all e-mail attachments and downloads before they’re opened. If there are settings for scans of ActiveX controls and Java for harmful content, use them.

Be sure to allow the program to create “clean boot” or “rescue” disks; you never know when you might need them to help you heal an infected system.

Chapter 5

Patching and Updating Your System and Software

In This Chapter

- ▶ Keeping up with patches, updates, and service packs
 - ▶ Preventing rootkits with regular updates
 - ▶ Getting the goods from Redmond — Windows Update and Microsoft Update
 - ▶ Knowing when you need a new system
 - ▶ Patching and updating application software
-

Current software patches and updates are critical to resisting malware attacks, especially if you're running Windows. No way around it — you need these. Malware and rootkit authors use the problems that the patches and updates repair as ways to compromise those hapless computers left unpatched and un-updated. Users who delay (or resist) getting the patches and updates may suddenly discover — way too late — that their computers are infected.

In this chapter, we discuss the importance of updating and patching your computer hardware, your operating system, and your software applications. This is a vital step in the prevention of rootkits. Just as you don't want to be caught out in the winter wind with holes in your jacket, you don't want to be caught out on the Internet with holes in your computer system.

Preventing Rootkits by Patching Your Clothes

Your operating system, software, and the programs you use are like the clothes you wear when you go outside. You want them to be clean, neat, and whole, depending on your personal preferences. Although some may consider it fashionable to appear dressed in tattered rags, the majority are well dressed and looking good. Cold winds blow most keenly through holes and tears in fabrics. What good is high fashion when you're freezing?

Okay, hold that metaphor: Holes and rips can appear in your operating system and applications, due to the wear and tear of everyday use or by attacks from nefarious people. Just as fabric can have weak spots or places where the seams come undone, so it is with software and operating systems. Human beings don't make perfect things; the manufacturer may have left unexpected flaws in the design and construction of your software. As we use the products, these design flaws eventually appear — and need to be fixed.

In computer lingo, *patches* are little pieces of software (and sometimes limited- or one-time use applications in themselves) that usually apply to design flaws in software, while *updates* include both patches and improvements to the product as a whole.

A big part of what criminal hackers are after is to search for weaknesses in your software and operating system so they can gain access. If they find holes and design flaws, they can easily get in and use your computer. They find these flaws by first working on test systems of their own, then using what they learn to apply it to computers on the Internet. They may share their findings with their criminal hacker buddies or may already be part of a hacking team. Prominent among the tools that criminal hackers use are rootkits — but before they can install one, they need to enter your computer. Finding an unpatched hole or design flaw is their ticket. They start by installing a backdoor so they can maintain access to your computer without having to log on. With the backdoor installed, they can then install the rootkit — which hides their presence on your computer.



Remember to check with all your software vendors regularly for patches and updates. Get the patches, updates, and Service Packs as soon as they're made available, installing them promptly or automatically; otherwise you could be faced with a stranger inside your computer. These fixes also apply to keeping viruses, worms, trojans, and many forms of malware off your computer.

Updating Your Operating System

Updating your operating system is an essential part of preventing rootkits. Exploiting unpatched security holes is probably the most common method of hijacking someone's computer, which is why patching known Windows security vulnerabilities on *patch Tuesday* (the second Tuesday of every month) is a must.

Get Windows Security Updates for your system from Microsoft regularly. Although some folks may disagree with this practice, consider: Without the security updates, your system is vulnerable to those problems so widely publicized by the media — after all, cyber-criminals read the news too. The Microsoft Windows Service Pack Road Map offers more information:

www.microsoft.com/windows/lifecycle/servicepacks.aspx



Vendors of application software will cease to provide programs and updates for old or obsolete systems. Without security updates, people using unsupported systems are more likely to be infected by malware. Microsoft announces when they consider specific versions of their operating systems obsolete — and then they stop supporting those versions with updates and patches. For example, Windows 98, 98 Second Edition, and the Millennium Edition all became obsolete June 30, 2006. Microsoft extended the date on those versions so customers would have time to upgrade their systems. Result: If you're a Windows user and want to be secure online, you have to get Windows 2000 SP4, Windows XP SP2, or Windows Vista if you haven't already. To check where your system stands, see Microsoft's Life Cycle Policy here:

www.microsoft.com/windows/lifecycle/default.aspx



Microsoft Update only supports Windows 2000 SP4, Windows XP SP2, and Windows Vista. Windows Update provides limited services to all unsupported versions of Microsoft Windows operating systems, primarily with high-priority security updates.

Patching, updating, and Service Packing

Microsoft regularly posts updates and patches for their operating systems and software. The updates or Service Packs contain improvements that can make your computer purr like a kitten, running more smoothly while increasing (and ensuring better) security. They offer these improvements for two reasons:

- ✔ They continue developing improvements for your system (technology marches on), partly as an incentive for new users to buy Windows.
- ✔ They may need to ensure your security when proven weaknesses or vulnerabilities crop up, found by Microsoft itself or by others.

These are widely publicized in the media — so the news becomes public knowledge in a hurry. So it can be of great importance to apply the service pack or patch to your operating system as quickly as you get wind of the need to do so. You may be given these patches as automatic critical updates. For more security and improvements than that, go to the Microsoft Downloads Center as described earlier in this chapter.



Hackers also read the news; they're counting on people to neglect to install updates and patches. Slack security makes their jobs *much* easier.

Looking at why you need updates

All updates and service packs are, technically patches — except they're typically very large and comprehensive (especially service packs). As we mentioned earlier in the chapter, a patch is usually a bit of software that fixes a specific problem with an application, most often bugs, or a weakness in the program. Updates and *service packs* often consist of a package that brings together many patches designed to eliminate numerous problems, improve the programming, and maintain your security.

The first month after a new patch is released is a critical period. Within the first week of a new update or patches, fresh malware will appear that uses the weaknesses the fix was designed to counteract. This is the time when unpatched systems are most vulnerable.



Here's one way they do it: Malware writers reverse-engineer a patch as soon as it becomes publicly available — and then write code to exploit the vulnerability it fixes. It was recently reported that it takes an average of six days for the patch exploits to start appearing on the Web after they were released. Blackhat hackers take full advantage of this window of opportunity to bring down any systems that haven't been updated in a timely manner. Your computer or servers are extremely vulnerable during this period; don't make some cyber-crook's job too easy.



Comparing 16-bit, 32-bit, and 64-bit architectures

You may be seeing references online now to 64-bit architectures. They are the next generation of operating systems and software. Architectures of the past were 16-bit (early Windows and DOS), both 16-bit and 32-bit (Windows 3.0, 95, 98, and ME), 32-bit (Windows XP and 2003 Servers), and now 64-bit capabilities. The 64-bit processors in Windows XP Professional x64 systems can still run 32-bit programs — a good idea, since so many of them are still out there. The number of bits represents how much data can be handled once by the CPU. The bigger the number, the more memory you can use with your applications and processing. 64-bit processors provide a huge amount of memory, far more than double the capacity of 32-bit. With 32 bits, the processor can access up to 4 GB of RAM — but with

64 bits you can do 16 EB. Any application that needs large amounts of memory will benefit — for example, 3D graphics and animation, digital content, games, photos, video, and audio to name only a few. Next time you upgrade your system, you may want to go 64-bit.

To give you a better understanding of what 16 EB means above: An EB is an *exabyte*, which is 1,152,921,504,606,846,976 bytes (or 1,000 GB). The range of bytes is kilobyte (KB), megabyte (MB), gigabyte (GB), terabyte (TB), petabyte (PB), exabyte (EB), zettabyte (ZB), and yottabyte (YB). All the printed material on our planet has been estimated to be approximately 5 exabytes, so 16EB is a whole lot of active memory in the 64-bit computer!

Knowing where you can get them

When Microsoft announces that you need a particular patch, update, or Service Pack, be sure to download and install it — because it's a sure bet that an exploit is already out there, or will be soon. For more information on what to do, visit this Web site:

<http://support.microsoft.com/kb/311047>

Taking advantage of Automatic Updates

Microsoft provides the means of acquiring individual patches and updates, either manually or automatically. Windows Update or Microsoft Update; manual or automatic; which should you choose? This section walks you through both methods.

Taking advantage of Automatic Updates is the safest way to make sure you receive Windows Updates as soon as they become publicly available. You can choose to do manual updates but please be sure to do them.)

If you have Automatic Updates enabled, then your updates download to your computer when you're online: A yellow shield icon appears in your system tray (in the lower-right corner of your desktop) while this is going on, showing a bubble that tells you Windows Updates are ready to download and install. When you see it, save anything you need, close any programs you have open, and then click the shield to begin the update. You can also let it run in the background while you continue your work. The downloaded updates are fully installed the next time you reboot. If you choose to install your updates later, the updater may prompt you repeatedly until you *do* restart or reboot your computer.



If you're not an administrator and the updater requires a reboot, it will reboot regardless — which is why we suggest saving your work first.

To enable Automatic Updates click Start ⇨ Control Panel ⇨ Automatic Updates. Or, you can click Start ⇨ Control Panel ⇨ System and click the Automatic Updates tab in the System Properties window. Many means exist to open the System Properties window and these are just a couple. Select the green shield at the top and set the time you would prefer to get them.

Guide to Windows Update and Microsoft Update

We emphasize repeatedly the importance of getting your updates to Windows and other Microsoft products in a timely manner. Windows Update has been the mainstay of all updates to the Microsoft operating system for many years. Most PC users are familiar with this program, at least to some extent. Windows Update helps you replace old and outdated system files with new ones — the basics — while also providing a means of installing needed patches for security threats.

The new kid on Microsoft's block is Microsoft Update. It includes Windows Update — while providing updates and patches for Microsoft products such as Office 2003 (and higher) and their other applications on Windows 2000 SP4, XP, and later operating systems.

Some confusion has occurred between using Windows Update and Microsoft Update. Windows Update is old and will be phased out eventually. Microsoft Update includes Windows Update and works only for currently supported

Microsoft operating systems (see “Updating Your Operating System” earlier in this chapter). High-priority security updates are still available for older, unsupported Microsoft systems from Windows Update only.



Windows server administrators can access server updates at Microsoft Update (as well as Windows Update).

The following sections give you the lowdown on using Windows Update and Microsoft Update.

Getting ready to use Windows Update or Microsoft Update

To use Windows Update or Microsoft Update properly on a Windows XP computer, here's what you need:



✓ To be logged on as an administrator. Automatic updates will come and install no matter who is logged on to the system, but to make changes to the updates procedures or manually download and install them you need to be an administrator or a member of that group.

✓ To be using Internet Explorer to access the Windows Update Web site.

✓ To have Automatic Updates enabled so you can download from the Windows Update site.



✓ A legitimate copy of Windows XP. If you discover that your operating system is not legitimate, Windows Update and Microsoft Update won't work, but you'll receive high-priority security patches and updates if you have at least Windows XP Service Pack 1 (SP1) installed and Automatic updates enabled. Some options exist by which you can obtain legitimacy at lower cost than if you buy a new operating system CD. Contact Microsoft for more help and details.

Using Windows Update and Microsoft Update

The easiest way to access Windows Update is to open Internet Explorer and then choose Tools ⇨ Windows Update. You can also get to it by clicking Start and choosing All Programs ⇨ Windows Update.



If you've upgraded to Microsoft Update, Windows Update won't appear anymore when you open Windows Update; that site will be bypassed and Microsoft Update will appear instead. The two services have similar but different interfaces, so the instructions for Windows Update may be different than with Microsoft Update.

At first, the Windows Update scans your computer to see whether you can use the site. When it's done with its inspection, if you can use it, the Web page will appear.

When you go to the Windows Update (or Microsoft Update) site, after the system scans to see what version of Windows Update or Microsoft Update software your Windows has, you'll see two buttons on the right pane:

- ✔ **Express (recommended):** Will provide high-priority security patches and updates suited to your particular individual computer. This is fine if you're a home user, but what if you're part of a network? Some of these security updates may be overridden by the network's configuration and settings.
- ✔ **Custom:** You choose which security updates and patches to install on your computer from both high-priority and optional lists, but again, these are suited to the individual home user.

Even if your automatic updates are enabled, there's more at Windows Update than that.

If you get an error message instead, follow the instructions provided to enable access again.

Your results will include updates in three major categories:

- ✔ **High-Priority Updates:** These consist of critical security updates and patches essential for the safety and well-being of you and your computer. If you have Automatic Updates enabled, you'll get these right away.
- ✔ **Hardware Updates:** These can be obtained from the Windows Update Catalog section. See the "Obtaining updates for networks and hardware" section later in this chapter for detailed instructions.
- ✔ **Software Updates:** These can also be obtained from the Windows Update site, the Microsoft Downloads Center, or Windows Marketplace.



When you make use of either update system, be sure to reset your ActiveX controls in Internet Explorer to Prompt (see Chapter 3) so you can decide manually, on a case-by-case basis, whether to allow ActiveX controls to run — especially if you want automatic updates. If you restrict ActiveX components by setting them to Prompt, then they will always ask before trying to download to your computer. If you choose to block ActiveX, you'll still have to allow it for this installation — and then update manually, again allowing ActiveX for the updates each time, and redoing the block when you're done. (See Chapter 4 for more information about ActiveX controls and security settings in Internet Explorer.)

This is where it all happens. Get your Automatic and critical updates first, then come back to see the Microsoft Downloads Center. Everything you need or want to try is available here, including updates for your other Microsoft applications, drivers and hardware devices. Here you can download Internet Explorer 7 and Windows Defender (both discussed in Chapter 4).

On the left pane, click Change Settings. Scroll down on the right of the next page. Put a check mark in the box beside Advanced and then click the Apply Changes Now button.

That brings you to the page at the Microsoft Downloads Center. If you have any problems opening this page, use this address:

```
www.microsoft.com/downloads/Search.aspx?displaylang=en
```



Don't go wild when you start downloading new enhancements and products. Set a fresh restore point first (see Chapter 3), and then download them one at a time. After installation, give them a test drive to ensure your system's stability and well-being. If you encounter problems, you can uninstall the applications and restore your computer.



You'll find lots of info by scrolling down this page, including links to each of the available download categories. Be sure to check your system requirements before downloading any applications.

In either case, if your computer is part of a network, you'll need to check with your network system administrator before getting updates.

Installing Microsoft Update

When you install the Microsoft Update application software and get your priority updates, when you go to Microsoft Update using the Internet Explorer browser, the service will scan your computer to see what updates you need and then present you with a list. These are updates beyond those which arrive via Automatic Update; they include updates for your Microsoft application software, such as Microsoft Office.

If you're still using Windows Update, you'll see an ad for Microsoft Update at the Windows Update site. Click the ad and follow the instructions to install Microsoft Update.



You will not see the ad if you have the Microsoft Updating program installed, or are using an unsupported operating system, such as Windows 98 Second Edition or Windows ME.



Note that Microsoft Update has been known not to work properly on every computer. After you upgrade to Microsoft Update, if it doesn't work properly (or if you just don't like it), you can get back to Windows Update by clicking Change Settings in the left pane of the Microsoft Update page, scrolling to the bottom of the page, and then checking the Disable Microsoft Update Software and Let Me Use Windows Update only option.

The Microsoft Updating program runs in much the same way as automatic updates do, in the background. You need to do nothing more than allow it to proceed.

If you do not have Microsoft Update installed on your computer, and do not have automatic updates enabled, you can access the Microsoft Update pages in order to install it. You need Internet Explorer 5.5 or higher to access the site. To install Microsoft Update, follow these steps:

1. To access Microsoft Update, go to the following URL:

```
http://update.microsoft.com/microsoftupdate
```

2. Click the Start Now button and then click Continue.

3. Read and follow the procedure to install an ActiveX control.

4. Click Install.

Depending on your computer configuration you may need to repeat Steps 2 and 3.

5. If you have Automatic updates enabled, go to Step 5.

- a) If not, click Turn It On Now and follow the procedure.
- b) Click Change Setting to turn Automatic updates on, following along with the procedure.

6. Click Check for Updates.

After your computer has been scanned, you are presented with the Microsoft Update Welcome page. See the "Using Windows Update and Microsoft Update" section earlier in the chapter for information on using Microsoft Update.

Obtaining updates for networks and hardware

If you're a network system administrator, you can obtain updates for multiple computers using Windows Update Catalog section of the Windows Update site. Home users too can find hardware driver updates via the method described here.

You need to be an administrator or a member of the administrator group on your computer to follow this procedure. You also need the Internet Explorer browser and either Windows Update or Microsoft Update (software) installed. If you haven't installed Windows Update or Microsoft Update already, you'll be prompted to do so.



If you intend to make any downloads from Windows Update Catalog, you'll have to temporarily disable the Popup Blocker in Internet Explorer. If it's left on, the downloads will fail. Here's how to turn it off:

1. **Choose Tools ⇨ Internet Options and select the Security tab.**
2. **Select the Internet Zone, then click the Custom Level button in the “Security level for this zone” box near the bottom of the Internet Options dialog window.**
3. **In the Security Settings box, scroll down until you see Use Popup Blocker (it's near the bottom). Select Disable and then click OK.**

If you get a Warning pop-up that asks whether you want to change the security settings, click Yes and then click OK twice.

When you've put Popup Blocker temporarily out of commission, follow this next set of steps to download Windows Updates:

1. **At the main window for the Windows Update site, under Options in the left pane, click Use Administrator Options.**

The following Web page appears. Network administrators will find many updating options and links at this page.

Although this procedure is not specifically dedicated to the needs of network administrators, it is comparable. If you're new to this process, it can help you.



2. **Click the Windows Update Catalog hyperlink (in blue text) in the Update Multiple Operating Systems heading at the top.**
3. **Click Find Driver Updates for Hardware Devices on the right pane, or Find Hardware Driver Updates on the left.**
4. **Click a category; for this example, you can click Video.**

This brings you to the Search function.

5. **Find and select the hardware manufacturer (in this case, that of your video card).**



If you don't know this information, from your desktop or the Start button, right-click My Computer, and choose Properties. The System Properties box opens. Select the Hardware tab and then click the button for Device Manager. This lists all the hardware installed on your computer. For example, you can click Display to get the name and model of your video card.

6. Select the operating system and the language and then click Search.

The results appear.



7. In the box at the upper-right of the Results page, select whether to sort the list by title or date.

Choose by date if you want the most recent update.

8. Find the update you want and click it.

A pop-up box appears, listing the names of the hardware you want to update.

9. Click the Add button to place each update in your Download Basket.

10. Click the Browse button and select a folder in which to place the download on your computer.

Please note that the path, including filename, must be a string of fifty characters or fewer. If it has too many characters, a pop-up will inform you.



11. Click the Download Now button.

You can also leave this page and find other downloads to place in your basket if you want.

You'll be presented with the EULA (End User License Agreement) for the download.

12. Read the EULA and click Accept.

A window appears, showing the progress of your download. When the download is complete, you're presented with a download history page.

For more information on Windows Update and Microsoft Update. . .

For more information on the many forms of Windows and Microsoft Update, see the following URL:

```
www.microsoft.com/downloads/render.aspx?displaylang=en
&content=updateservices
```

For a FAQ on Microsoft Update, see the following URL:

```
www.microsoft.com/athome/security/protect/update.msp
```

For instructions on how to enable Microsoft Update to run automatically for Microsoft Office or other Microsoft programs, please refer to the following Web site:

www.microsoft.com/athome/security/update/msupdate_keep_current.msp

Patching and Updating Your Software

Most vendors offer the capability to update their software manually or allow the software to search periodically for updates on its own. This section reviews the methods you can use to obtain patches and updates for your Microsoft compatible applications, points you toward some online help, and offers pointers on matching your updates to the way you use your system.

Some Microsoft-compatible applications you use may not need updating, but if you're using vendor-supplied anti-malware programs, they *do* need updating and patches — on a regular basis. Some of these, such as antivirus programs, may need their signature files updated every day. These programs usually come supplied with the means of obtaining automatic updates so you can set it and forget it. But what if you don't like automatic updates and would rather do this manually? We discuss how to do that in the next section.

Ways to patch or update your applications

Just like your operating system, your chosen applications may need updates and patches depending on their purpose. Vendors of anti-spyware, antivirus, anti-trojan, firewall, and other security programs provide these on a regular basis — sometimes even daily. Check your applications' help files to learn how to configure them for automatic or manual updates.



- ✓ Many vendors — especially those who provide security applications — also provide Web sites where you can read about the latest patches, fixes, and updates for their software. Often the applications feature links to these sites in their help files or even right on the program's Help menu. Other online resources include official forums where you can ask questions and get help.
- ✓ A simple way to find forums for your applications — in the language of your choice — is to search in Google with their names. Open a blank page in your Web browser, type **www.google.com** into the address bar and press Enter. Then, on the Google page, type the name of your software and click the Search button.

Watching Internet sources for known problems with your applications

Another way to learn of updates and patches is to check a Web site dedicated to providing information on a great many of them. Calendar of Updates is such a site; crewed by a group of dedicated volunteers, it tells you what applications are being updated when:

```
www.dozleng.com/updates/index.php
```

Yet another is the Updates Forums at the CastleCops Web site:

```
www.castlecops.com/c13-Updates.html
```

You can learn more about new application vulnerabilities at these sites:

```
secunia.com
```

```
www.securityfocus.com/vulnerabilities
```

Patching and updating shared computers in heavy use

Systems that have multiple users are a little harder to keep patched and secure. You may have a system you share with your spouse, roommate, or co-worker. You may be in charge of systems that multiple users work on. Or you may need to use a system that you have no control over, such as one at an Internet café or airport. What should you do?



You can use this book's CD to run some of the online scans listed on the CD from these systems to check for basic malware — but clean results are no guarantee of security.

When you share systems with other users, or you're the administrator of these multiuser systems, it's a good idea to increase the intensity of your housekeeping routine. You must be the responsible one; don't rely on the other users to keep the system up to date. Better to keep up than have to catch up.

In all cases, you should have your anti-malware software doing daily automatic checks for updates — if not more often. All the anti-malware scanning software should remain resident, which means the real-time monitoring functions are on at all times.



Here are some pointers for getting the most rootkit protection out of your anti-malware program:

- ✓ Most anti-malware software has a password (which you can enable) that's required to disable the "auto-protect" feature; people who don't have it can't just turn off the software if it annoys them.
- ✓ Have your anti-malware software run checks at user logon. This may seem to slow down the users' effective work time — but think of it as an investment: The time it takes to clean an infected system, plus the time taken up by the damage control that follows, will most likely be a far greater drain on productivity than the ten-or-so minutes it takes to run those scans, even if you run them two or three times a day.
- ✓ If you're a home user who shares your computer with others, you might want to create a fresh restore point before and after each time you use the system.

Spend a little time to save a lot of time; it's that simple. The time spent in proactive maintenance will far outweigh the time required for reactive repairs.

Knowing When You Need a New Computer

Okay, this may hurt your wallet a bit. But it's an inevitable reality of working with computers: system upgrades. A decade ago, you had to get new hardware every couple of years in order to use the newest and latest application software. Today (fortunately), the hardware has surpassed the software — so you can keep using your computer for up to five years, perhaps even longer. But sooner or later, we all have to pay the technological piper.

If your hardware is more than three years old, it may not be able to run Windows XP — and don't be surprised if what you're using now is completely flummoxed by Windows Vista. A telltale sign is also when you try to run a new application and your older hardware does not have the power (in particular the CPU speed or sufficient RAM) to run it. Therefore, in addition to upgrading your operating system, you may need new hardware. Check with your vendor or local computer dealer for more details. Smart shoppers check local computer-user clubs and look around online to find out what's hot and what's not. The computer-user clubs can be found in your local phone directory (it's much quicker than searching for 'em online). Tell them what you want to get or what you think you need, and ask them what they'd suggest. They'll likely give you a better outlook on it than you could get in a sales situation.

Chapter 6

Blurring the Lines of Network Security

In this Chapter

- ▶ Understanding and using Windows access-control mechanisms
 - ▶ Limiting and controlling physical access
 - ▶ Controlling removable devices
 - ▶ Setting up installation barriers
-

In the past, server networks had different security needs from those of standalone computers. No more; the threat posed by rootkits has changed all that. Network security is still the top priority for the broader base of computer users, but the workstation can't be viewed as secure if it just sits there behind the network firewall. Workstations have to be secured in much the same way as standalone computers — rootkits (among other threats) call for this essential layer of additional protection against network compromise.

Depending on your particular setting, whether at home or at work, you have some practical means available for developing your computer security to a higher standard. This chapter describes some of them — and lays out what you need to do to begin developing stronger security, both online and off.

To ensure a safe and progressive future for your home and business — or at least your computing capability in either place — you have to release your inner security geek. Running standard security programs — say, a firewall and an antivirus — is still important, but that's not nearly enough to cover your computer's assets in today's online environment. Here are a couple of good starting points for improving security:

- ✓ **Hardening your Windows operating system:** This process involves setting better security standards, keeping your patches up to date, and using applied programs to fill the gaps in your defensive arsenal. These tasks are discussed in detail in earlier chapters in this part.

- ✔ **Understanding how this security stuff works:** Here the idea is to get a handle on monitoring and maintaining the vulnerable parts of your system. Don't expect a "guru" to do it for you, unless you have money to burn. Although the intricacies of data processes may be beyond our grasp, we can discover ways to apply what we do know to make things better and more secure.

A Checklist for Improving Security

Half the fun of learning how to secure your computer for Internet use is simply discovering what you need most. Here's a no-frills, basic recipe for cooking up better protection for your network servers, workstations, and standalone computers:

- ✔ Configure your computer using customized Security Templates. Note that the security configuration of the network will override local settings on workstations. See the "Using Windows Access Control" section later in this chapter for more information.
- ✔ Enable auditing of computer events, whether or not on shared equipment. (The "Learning to love auditing" section later in this chapter provides more information.)
- ✔ Set up Limited Access accounts for online use, and require strong passwords to log on. (See Chapter 4 for more information.)
- ✔ Be sure to get — and maintain — your Microsoft Security updates. One approach to seriously consider is setting the updates to install themselves automatically. Keep your software patched and updated as well. (Chapter 5 has the lowdown on this process.)
- ✔ Disable the Windows Firewall. Choose and install a more effective personal firewall, and configure it to high security. Even better, use a router or hardware firewall, in addition to a personal one — especially for network servers and standalone computers. Chapter 4 gets you started in this direction.
- ✔ Install an antivirus application. Keep it updated — we recommend setting it to do its thing automatically — and on a daily basis.



Be sure to do regular security checks at least twice per week: Go offline, temporarily disable any unneeded programs, and run full-system scans with your anti-malware programs. (See Chapter 4 for more information.)

- ✔ Install anti-trojan and anti-spyware applications. These work similarly to antivirus programs — and require similar care and feeding. Keep them updated and run full-system scans with them twice a week — or more often, depending on how much time you spend online. (Chapter 4 offers guidance.)

- ✓ Using encryption software can tighten up security and privacy for your network. You can enable file encryption in Windows XP Professional via the Encrypting File System (EFS) that comes with it. You can learn more about EFS by entering those three letters into the search box in your Help and Support section in the Start menu. You can also read more about it at the following URL:

www.practicalpc.co.uk/computing/windows/xpencrypt1.htm



For Network Administrators: You might want to brush up on RestrictAnonymous and the LSA key, to disable anonymous enumeration through null sessions to improve security. More detailed info is available from the following online article:

www.securityfocus.com/infocus/1352

Learning to Love Auditing

It's pretty common for ordinary words to take on specialized meanings when applied to computers. So what's an event — and why would you log it? In programming terms, an *event* is any action — normally initiated by a user, Windows, or running background programs — that requires a response from the system. Behind the scenes, your Windows XP or Windows Server 2003 computer silently keeps a history of all the significant events which take place on it. These events are recorded in three separate logs, each one based on the specific nature of the data it collects: Application, Security, and System. This process is called *auditing* or *event logging*, and can be compared to having a security guard watching everything that is going on in your operating system and recording anything unusual that occurs — according to your instructions — in a daily report. Sometimes malware betrays its own presence with telltale events — and if we're hunting rootkits, we're looking for events such as repetitive logon failures, new service and driver installations, or Stop errors produced from poorly coded rootkit drivers, just to name a few. (These signals get the detailed treatment in Chapter 8.)

The events collected in the three types of event logs can be viewed by using the Windows Event Viewer. They are stored in %systemroot%\WINDOWS\System32\Config folder with the suffix .evt on Windows XP (where %systemroot% represents the drive letter where Windows is installed, such as C:). The Event Viewer can greatly assist your security by helping you piece together how malware entered your computer or network. It can also help you gauge whether your operating system and applications are working properly (not every failed Windows event causes a system crash!).

In Windows XP Home Edition, all security events are recorded in the security event log by default. Windows XP Professional and Windows Server 2003 offer a greater degree of control by giving you the ability to customize which security actions you want monitored using the Local Policy Editor on a stand-alone computer or the Group Policy Editor for a network server. These operating systems also have auditing disabled by default (we show you how to enable auditing in the next section).

Whether you're a home user or a network administrator, it's a good idea to check all the event logs periodically, using the Event Viewer to keep tabs on what's going on. In this chapter, we concentrate mainly on how to enable and configure security auditing on Windows XP Pro and Windows Server 2003 — but Windows XP Home users should come along for the ride to get the low-down on some general event logging and security principles that are relevant to all Windows users.

Enabling security auditing

The first thing called for here is an attitude adjustment — because auditing is often overlooked. As mentioned already, on Microsoft Operating Systems, such as Windows XP, Windows 2000 Professional, and Windows Server 2003, security auditing is disabled by default. By enabling security auditing, you can monitor specific events on your computer, which may help identify how your system was compromised by malware such as a rootkit. This gives you a starting point for investigating theft or corruption of data.

After you've installed the programs that can detect malware in your system, your first step is to enable the security-auditing features of Windows to harvest the clues that malware may have dropped.



Two excellent resources to help you get a handle on setting up different aspects of auditing — and clarify the importance of those auditing Event IDs — are

- ✓ Microsoft TechNet: Events and Errors for Windows and the Windows Server System:

www.microsoft.com/technet/support/eventerrors.aspx

- ✓ Microsoft's Security Monitoring and Attack Detection Planning Guide:

www.microsoft.com/downloads/details.aspx?familyid=95A85136-F08F-4B20-942F-DC9CE56BCD1A&displaylang=en

Turning on event logging

Windows systems come with system and application event logging turned on by default — already installed as a *service* (that is, a support program that's available to all components of the system) and set to start automatically every time your system restarts. Security logging is enabled by default on Windows XP Home (and cannot be turned off) but is disabled by default on Windows XP Pro or Windows Server 2003. If you find that event logging has been switched off (as described in Steps 1 through 3 in the following list), you can turn it on again by following one of the upcoming procedures, whichever corresponds to your operating system.

You can turn on event logging in the Services console (a special utility that controls only services) by doing the following:

1. Click the Start button and choose Run.

The Run window appears.

2. Type `services.msc` into the Open: field and Click OK.

The Services console appears.

3. Scroll down the services listed alphabetically in the right pane to find Event Log.

- If the word `Started` appears to the right of Event Log (under the Status column), you know event-logging service is already started and can close the console.
- If the Status column is blank, double-click Event Log and continue to Step 4.

The Properties dialog box for Event Log appears.

4. Click Start in the Service Status field.

5. From the Startup Type drop-down box, choose the way you want Event Logging to start.

Here you have two choices:

- **Automatic:** Event logging starts at Windows startup. This is the best choice, and the Windows default setting, because it enables information to be gathered constantly and stored in logs existing on your system. You can consult these event logs to keep tabs on what's going on your computer or to aid you in system troubleshooting.
- **Manual:** Event logging starts only when you turn it on as described in the preceding steps.

6. Close the Services console.

If you're looking for a faster way to start the Event Log service, you can do it directly from the command line:

1. Click Start, choose Run, type cmd into the Open: field, and press Enter.

A window (known as the *command console*) appears, displaying a command prompt.

2. Type `sc start EventLog` at the command prompt and then press Enter.

If the service is already started, you'll get a message to that effect, otherwise you get a response indicating the service is starting.

3. Close the command console.

These steps turn on event logging for all event types in Windows XP Home and Pro. If, however, you're running Windows XP Professional or Windows 2003 Server, you need to specify exactly what security events you want to audit by turning on security auditing. This selective monitoring can help you zero in on the actions that indicate whether your computer or network's security has been breached. Careful monitoring can also help you prevent an attack. The next section describes how to enable and configure security auditing to catch and trace threats that may compromise your system.

Turning on security auditing (Windows XP Pro and 2003 Server)

Windows XP Professional and Windows 2003 Server come with security auditing of all events turned off by default. Call it a gentle hint: You'll have to configure this feature to fit your system — not only turn it on, but also tweak it to monitor the success or failure of up to nine different security events. We list those events here — we figured they'd come in handy — basing our descriptions on categories found in the online Microsoft TechNet Document. You can find that useful resource online at this address:

```
www.microsoft.com/technet/security/topics/auditingandmonitoring/securitymonitoring/smpgch02.mspx
```

Fortunately, turning on and configuring security auditing in Windows XP Professional is the same procedure as for Windows Server 2003, with one exception: In Windows Server 2003, you can configure security auditing differently for each individual user. You can do that by specifying the Audit Privilege Use policy setting, as follows:



The following instructions are for Windows 2003 Server only.

1. Click Start, choose Control Panel, and double-click Administrative Tools (or, in Category View, click Performance and Maintenance and then click Administrative Tools).

2. **Double-click Local Security Policy.**
3. **Expand Local Policies and highlight Audit Policy.**
4. **In the Details pane, double-click Audit Privilege Use.**
5. **Place a check beside Success and/or Failure and click OK.**

Now, whenever a user exercises a user right, it's logged according to the choices you make in these steps. The audit settings can be applied to each individual user. You can check these in the Security log by using the Event Viewer, as described in Chapter 8.



You must be an Administrator or Administrative User to turn on and configure security event auditing. (Can't have just anybody mucking around in the vitals of the system.)

Here's the easy way to access the Security Policy editor to adjust your local security settings:

1. **Click Start ⇨ Run.**
The Run window appears.
2. **Type secpol.msc into the Open: field.**
3. **Click OK.**

Told you it was easy.

You also access the Security Policy editor by using a pretty straightforward series of clicks in the following steps:

1. **Click the Start button and choose Control Panel.**
The Windows Control Panel opens.
2. **Double-click Administrative Tools (if your Control Panel is in Category View, click Performance and Maintenance and then click Administrative Tools), double-click Local Security Policy, and then open Security Settings ⇨ Local Policies ⇨ Audit Policy.**

This will display the settings that define your current audit policy. You may adjust the policy settings in the details pane on the right side of the window by checking those events you elect to monitor. For each of the nine audit policies listed (and described in the next section), select to do one the following:

- Not audit
- Audit successful events

- Audit failed events
- Audit both successful and failed events

3. When you've made your choices, click Apply and then OK to confirm your changes.

Figure 6-1 illustrates our suggested security-audit settings.



Be sure to leave Object Access and Privilege Use set to Failure, or your log will fill up very quickly with routine, everyday stuff — making it harder to spot events triggered by malware.

The security auditing settings illustrated in Figure 6-1 enable you to use the Log Parser SQL scripts listed in Chapter 8 (and posted in the CastleCops Rootkit Revelations Forum at the address below). You can adjust the settings to conform to any auditing policy you've established.

www.castlecops.com/f233-Rootkit_Revelations.html




Figure 6-1:
Audit-configuration settings — these work for us.

Categories of security events you can audit

Here's a basic list of the types of event you can audit in your quest for hidden malware, with a brief description of how the auditing service does its job in each case:

- **Account logon events:** Attempts to log on to a local user account — an event that will appear in the log. If the user account is a domain account, then this event also appears on the domain controller.



A *domain* consists of a group of computers centrally managed on a network with mutual rules and procedures. Member computers have domain accounts by which to logon to their particular group. The *domain controller* is a server which stores member domain accounts information, checks members credentials, and maintains security for the group.

- ✓ **Account management:** Audits anything to do with creating, modifying, or deleting user and group accounts. It also keeps track of changes to passwords or password resets. Rootkits can create a new user administrative account to gain access to the system or use a legitimate account (by stealing passwords) to gain control of administrative privileges.
- ✓ **Directory service access:** Attempts to access objects in the Active Directory service. Rootkits affect the Active Directory service by adding new system resources to the Active Directory. The Active Directory pools all resources in a centralized location for the benefit of networked users. A rootkit can alter or add to the content present in the active directory and manipulate it to its own advantage.
- ✓ **Logon events:** Attempts to log on to computers, including workstations and member servers. Repeat unsuccessful logon events (logon failure) can signal that an anonymous user is attempting to access the system. In other words, it can be symptomatic of an intrusion or break-in attempt.
- ✓ **Object access:** Attempts to access an object such as a file, folder, Registry key, or printer that has audit settings defined in its system-access control list (SACL). Rootkits add information to the Registry to ensure their survival and install new files to gain control and achieve their goals (system compromise).
- ✓ **Policy change:** Audits any change to assigned user rights or to policies governing audits, accounts, or levels of trust granted to particular users.
- ✓ **Privilege use:** Audits each instance when a user exercises an assigned right, such as changing the system time. A rootkit needs administrative privileges to operate and it can even upgrade or modify security objects (tokens) to achieve its objectives. The rootkit known as Shadow Walker upgrades process privileges and its bag of tricks will be discussed in detail Chapter 12.
- ✓ **Process tracking:** Audits application behavior such as program starts or terminations. Rootkits and malware in general, will often try to disable security programs so their work can proceed without interference. They will also create new hidden services and processes, or use existing services and processes to execute and/or gain remote access.
- ✓ **System events:** Audits computer-system events such as startup and shutdown, and events that affect system security — or, for that matter, the security log. Examining system events is crucial to investigating clues to rootkit behavior. Kernel-mode rootkits achieve their effect by installing new services, and this will be logged as a system event.



A rootkit may conceal its own existence by interfering with the logging process, even to the extent that nothing significant will be reported in the event log other than evidence of the initial break-in attempt. That evidence will remain, because the intrusion evidence existed before the rootkit became active and put measures in place to hide its tracks. This, however, represents an extreme case of rootkit trickery, and often the event log will provide ample clues to a rootkit's existence. At the very least, the logon events should reveal some indication of an intruder's attempts to gain system access with the ultimate goal of system compromise.

Spotting suspicious activity using Event Viewer

After you enable the event log to record audits for all processes that stop and start, and for all system events, you have a foundation for using the Event Viewer application to hunt for rootkit activity — or for the doings of the malware that rootkits aid and abet. Chapter 8 provides the details of how to view, interpret, and respond to the information you get from Event Viewer as you're looking for clues the bad guys may have left behind.

Using Windows Access Control

The following utilities are known collectively as *Windows Access Control Mechanisms*: the Local Security Policy Editor, Security Configuration Manager, and the Microsoft Management Console. These tools together allow you to give specific permissions to users, groups, and computers for particular actions on a network or on a standalone machine. Along with auditing and event logging, it's like having your own security surveillance service for your computer or network. Files, folders, programs and Registry keys are all considered to be the objects that can be read, changed, saved or deleted. These tools help you to decide who, and what has the permissions to modify those objects on your system. The following sections show you how to use these tools to secure your computer.

Editing policies and configuring security

The specific features you use to secure your Windows operating system can vary by version. In Windows XP Professional, the Local Security Policy Editor and Security Configuration Manager are what you use to make or modify the security settings on your local computer or on your network. They provide you with a range of security-oriented tools and tasks:

- ✔ **Security templates:** Each template is a set of security instructions you can give your computer — stored as a file you can apply or remove as required. You can use templates for the local computer or apply them more widely to a Group Policy on a network.
- ✔ **Security configuration and analysis:** Security Analysis allows you to check to see whether the security levels you've chosen for individual computers or for an entire network are as you need them to be. It provides suggestions on how you can make corrections or adjustments with visual flags. You can make any changes you require with it.
- ✔ **Security configuration:** Use these settings to make direct changes to security on your local computer system — for example, by importing and applying security templates.
- ✔ **Security settings:** You can use an individual computer on a network to edit the security settings for local computers or for a Group Policy object that affects an entire group on the network.
- ✔ **Local Security Policy Editor:** You can use this application to establish or make changes to specific accounts — and to local security policies on your individual or standalone computer. You can control access to your computer, monitor the resources used, and audit or log the activities of users.

Making your own security-analysis utility

Okay, the do-it-yourself approach may seem a bit chancy at first, but you can put together a security-analysis utility for Windows to help you fight malware and rootkits. Building that tool is surprisingly straightforward; this section guides you through your first home-made security test. The skills and procedures you learn here will help you later on when you start working with the Local and Group Policy Editors. You'll need to be logged on as an administrator, or as a member of that group to use these techniques.



The following instructions apply to Windows XP Professional, and aren't available for the Windows XP Home Edition.

Testing your system against a security template

To get started, you have to get an accurate picture how your present security stacks up. So the first order of business is to test your system against a security

template; your operating system comes with a number of them. The templates range in strength from comparatively weak (`setup security.inf`) to very strong (`hisecls.inf`). For the demonstration in this section, we chose a medium-security template, `securews.inf`, as the starting point. (You can get the details of this template by choosing Start ⇨ Help and Support, putting Security Templates in the search box, and clicking the green arrow.) Before continuing, check your System Restore points — or set a new one in case anything goes wrong. See Chapter 3 for more information.

Getting snappy with a snap-in

In this section we show you how to use the Microsoft Management Console (MMC) to compare your current security settings to those of a medium security template. You can add the template as a snap-in to the console. The MMC is a hierarchal framework that contains administrative tools for organizing, monitoring, and securing the local computer or network. The comparison gives you a handle on what you need to do to improve your security settings to prevent, observe, and warn against malware and rootkits.



Note that we're *not applying* this template yet — we're only testing the current security settings of a local computer against it. Don't apply this template to your computer unless you're already enough of a computer guru to know exactly what you're doing and what can go wrong. These instructions are for standalone or individual computers, they're not for computers in workgroups or those on local area networks. Even for this demo, you must be logged on as an administrator to proceed.



It's possible to apply the template to the computer in one fell swoop but it is *not* advisable to do so. The security templates are used for tests and comparisons but aren't normally to be applied across the board. It's a temptation to apply it to force all the old settings to change to new ones. It isn't so easy as that here. The template represents a large number of ideal and specialized security settings which can make your computer go crazy if applied all at once. You're supposed to change a few settings at a time and make sure everything's okay before doing a few more.

That said, here's the template test, step by step:

- 1. Choose Start ⇨ Run, type mmc into the Open: field of the Run window, and click OK.**



If the Run command isn't already enabled in your Start menu, right-click the Start button, and then choose Properties ⇨ Customize. Then select the Advanced tab, scroll down the Start Menu Items list until you see Run Command, and put a check mark in the box beside it. Click OK ⇨ Apply,

and click OK when the next box appears. The Run command will now appear in your Start Menu.

The Microsoft Management Console appears, as shown in Figure 6-2.



Figure 6-2:
The
Microsoft
Management
Console.

2. Choose File ⇨ Add/Remove Snap-in.

The Add/Remove Snap-in window appears, as shown in Figure 6-3.



Figure 6-3:
The Add/
Remove
Snap-in
window.

3. Click Add.

The Add/Remove Snap-in window appears, as shown in Figure 6-4.

4. Select Security Configuration and Analysis from the list, click Add, and click Close.

Figure 6-4:
The Add Standalone Snap-in window.

5. Click OK on the Add/Remove Snap-in window.

After you add the Security Configuration and Analysis Snap-in to the Console Root, your Console will look like Figure 6-5.




Figure 6-5:
Security Configuration and Analysis Snap-in added to Console Root.

Making a database of your settings

Now, in order to save your work and so you can test your system with the template, make a database of your settings. The database of your template settings is important to have; without it, you can't make any comparisons. Here's the sequence of steps to create it:

1. **Right-click Security Configuration and Analysis in the left pane, and then choose Open Database from the shortcut menu.**

The Open Database window appears.

2. **Type in a filename of your choice for your database and press Enter.**

The Import Template window appears, as shown in Figure 6-6.




Figure 6-6:
The Import
Template
window.

3. **Choose `securews.inf` from the list and click Open.**

Note that if your system is not set up to show file extensions, the template name appears simply as `securews`.


4. **Click the Security Configuration and Analysis snap-in in the left pane.**

Your console should look like Figure 6-7.

5. **Use File ⇨ Save-As to save your new security tool to the default Administrative Tools folder; name it whatever you want.**

Now you won't have to retrace these steps to rebuild the tool in the future — and you can open it easily from Administrative Tools in your Start menu.

Figure 6-7:
Console
Root\
Security
Con-
figura-
tion
and
Analysis.



Testing your system

The next stage of tightening up your security is to do a security check of your system that compares it to the template we've chosen here. Instructions on how to proceed should appear in the right pane, as in Figure 6-7. The steps go like this:


1. Right-click Security Configuration and Analysis in the left pane.

A context menu appears.

2. Choose Analyze Computer Now and then click OK to accept the default path for the log file.

After a few minutes, the results appear on-screen, looking something like Figure 6-8.

Figure 6-8:
Security
Analysis
Results.



3. Examine your results.

They appear in the Security Options window shown in Figure 6-9.

- If the icons have little padlocks on them, it means that these were not included in the security scan.
- If the icons don't have padlocks, click their + signs to see their sub-headings in the left pane.
- Click a subheading in the left pane to see its components in the right pane.



Figure 6-9:
Security
Options look
like this.

If you see a lot of green circles with check marks in them on the icons, then your system is already configured to a security level that matches the template. If (instead) you see a lot of red circles with X signs in them, then your system is not very secure as compared to the template — which is (uh, yeah) really handy to know.

The columns in the right pane show how your system's security configuration compares to the template we loaded for this demo. The Policy column consists of the component settings; the Database or Template settings and your actual Computer Settings are also shown.

This demo shows you what needs to be improved. Don't attempt to make changes here as they will only affect the database you've created and not your computer. You can make changes to the security settings for a stand-alone computer using the Local Policy Editor (see the next section).

Interpreting the results of your analysis

Okay, when you have results to work with, what do you do with them? You can choose one of three ways to deal with these listings:

- ✓ **Leave everything as it is.** If you're happy with your settings and (after reading this book, of course) confident that you're safe from malware and rootkits, then don't change anything. It's up to you.
- ✓ **Change any of the settings in the Policy column, one at a time, using the Local Security Policy Editor.** You can access this column, along with the instructions it provides, by following these steps:

1. **Click Start and choose Help and Support.**
2. **Type local security policy into the search box and click the green arrow.**
3. **Choose Pick a Task ⇨ Edit Local Security Settings and follow the instructions that appear.**

You can also access your local security settings from Administrative Tools in your Start Menu.

- ✓ **Apply the security template settings to your computer directly.**

If you're an advanced user, you might get away with this — but Microsoft warns against doing so. They assert that it can have unexpected effects. To apply the template, do the following:

1. **Be sure to set a fresh restore point before you begin.**
See Chapter 3 for instructions on using System Restore.
2. **Open the security tool created earlier in this chapter.**
Or, follow the procedures in the previous sections to create one.
3. **Right-click the Security Analysis and Configuration heading in the left pane and choose Configure Computer Now . . . from the context menu.**

Windows lets you know when it finishes making the changes. You may need to restart your computer for the changes to take effect.

Take your computer for a test drive to see if you like the new settings. If you don't, then restore your computer using the restore point you made in Step 1.



With the Local Security Policy Editor, only the settings you change are applied to your system.

Configuring security selectively

To make specific changes to your system's configuration using the Local Security Policy Editor, follow these steps:

- 1. Double-click a Policy component in the right pane.**

A dialog box specific to that configuration appears.

- 2. To change the setting you double-clicked, uncheck or check the boxes beside Success or Failure.**

See the “Learning to Love Auditing” section earlier in this chapter for more information.

- 3. Click Apply, and then click OK.**

Repeat these first three steps for any other settings that seem necessary to you — and when you’re done restart your computer.

- 4. Try out your Internet connection, e-mail, and any other applications to see how they are affected by the changes.**

If you have problems, you can use System Restore to get back to your last known good configuration — and start over.



To err on the side of caution, we suggest you change only a few settings at a time. Then restart and see how (or whether) they work before doing more.

Your Help and Support section can provide more info about the individual settings. If you want a bit wider perspective, try searching Google for terms such as *security templates*, *local security policy editor*, and *security and configuration analysis*.



These particular settings only apply to single, standalone computers. If your computer is part of a network, remember that network settings can override yours. If those network settings are less secure, then you’ll be back where you started as soon as the network refreshes its settings and “updates” your computer. However, you do need these settings in order to create a secure database for network applications. See the next section for more information.

Customizing a security template for a network

Securing your business network is a little more complicated and involved than doing a standalone computer, but the two processes have something in common. The previous section shows how to use the security template — and that’s the first step in securing a network of computers. After you’ve created a security database, you can apply it to more computers by importing it into a Group Policy and then applying it to a network. (You can use these procedures to configure the local policies of a computer from a custom template as well.)

If you want use one of the default security templates, you can apply it directly to your servers and workstations. Or you can change some of the settings to

customize it to meet your needs — but before you do that, make a copy of the default security template in case you make mistakes or change your mind. To make a copy of a default security template, do the following:

- 1. Add the Security Templates Snap-in to your Microsoft Management Console.**

The painless way to get this done is to follow the steps in the “Getting snappy with a snap-in” section earlier in this chapter, substituting the Security Templates snap-in for the Security Configuration and Analysis snap-in mentioned in that list. Figure 6-10 shows you what to look for.




Figure 6-10:
Security
Templates in
the Console
Root.

- 2. Under Security Templates, in the left pane find `C:\WINDOWS\security\templates`.**
- 3. Right-click the template you want to copy in the left pane, choose **Save As**, and enter a new name for it.**

The copy is saved in the same folder as the other templates. Don't change the filename extension.

With the new template, you can edit the settings manually by browsing through the headings in the Security Templates console (as described in the “Testing your system” section earlier in this chapter).

Your newly renamed security template appears in the headings list in the left pane of the console, as shown in Figure 6-11. Now you can drill down to the individual components to make changes.

- 4. Click the + beside headings and subheadings in the left pane until there are no more; then select the heading without a + sign.**

The heading's components appear in the right pane. In Figure 6-11, you see that Password Policy has been selected.



Figure 6-11:
Working
with a new
template.

5. Right-click a component in the right pane and select Properties.
6. Be sure there's a check in the box beside Define This Policy Setting In The Template.
7. Make your changes and then click OK.
8. When you're done making changes, save your work on this new template by clicking File ⇨ Save.

Copying settings from one template to another

You can also copy groups of settings from one template to another. To copy the settings from `securews.inf` to a blank template, for example, do the following:

1. Find `C:\Windows\security\templates > securews.inf` in the left pane.
2. Right-click the Account Policies heading and choose Copy.
3. Find the Account Policies heading in your new policy, right-click, and choose Paste.

Doing so adds the account policies from the `securews.inf` template, leaving the other headings unspecified.

Importing a security template into a Group Policy Object

After you've created a template with the security settings you want, you can import the template's settings into a Group Policy Object (GPO) so you can apply those changes to the computers on your network.



Network administrators please note: The settings in the Account Policies heading — account policies, account-lockout policies, and Kerberos policies — can be set only at the domain level in Windows.

Here's the drill, step by step:

- 1. Open the GPO to edit from the Group Policy MMC or open the Group Policy Editor by clicking Start and choosing Run, typing gpedit.msc into the Open: field, and clicking OK.**

The Group Policy window opens, as shown in Figure 6-12.

- 2. In the left pane, look under Computer Configuration and click the + sign beside Windows Settings to reveal Security Settings.**
- 3. Right-click the Security Settings heading and choose Import Policy.**




Figure 6-12:
The Group
Policy
window.

- 4. Select the security template you want to use from the list, and click Open.**

If this is the first time you've used Group Policy, you'll see a warning box that tells you that the Group Policy settings cannot be determined, that



the local security settings are not defined in Group Policy, and that network settings (domain-level policies) will override them.

When you import new security settings via this method, the Group Policy object does not accept it as a change (hence the Security Templates box). Fortunately, the next step takes care of that little detail.

5. **Change *one* of the settings in the GPO manually by clicking the + signs beside the headings in the left pane until the + signs are all gone. At this point, select the policy.**

The policy's components appear in the right pane.

6. **Right-click a component in the right pane and choose Properties.**

A box opens; here you can make that one manual change.

7. **Change the setting that you want; then click Apply and click OK.**

The idea here is to implement the changes automatically when the network refreshes the settings. (You can change the one you modified back to its previous setting later.) If you give it an inch, you will get a mile. That one little change will set up a condition that will cause all the changes to be accepted into your Group Policy.



The Group Policy settings you just applied are only on your standalone computer. You'll need another tool to apply it to your network — which, fortunately, we describe in the next section.

Applying Group Policy Objects to your network

The Microsoft Group Policy Management Console (GPMC) is what you use to apply Group Policy Objects to your network. The whole song and dance of doing that is far too lengthy to be included here. For more information and help, refer to the following link (which leads to the relevant Microsoft TechNet article):

www.microsoft.com/technet/security/smallbusiness/prodtech/windowsxp/netprtct.msp

Preventing Attacks by Limiting Access

To stop or catch a thief, you need to think like one. Do you see opportunities that blackhats could use to their advantage in your home, business, workplace — or in public? We've seen users leave their laptops open on their tables while they get a second cup of coffee. What would it cost you in terms of lost time, data, and money if your computer was stolen? In this section we

offer insights and practical means by which you can safeguard your equipment, your software, and your wallet.

Limiting and controlling physical access

To stay clear of rootkits, you have to protect not only your software, but your hardware as well. Locking things up physically may seem a little quaint in the twenty-first century, but criminal hackers and computer thieves can be anywhere. Businesspeople (for that matter, private individuals) with laptops are increasingly mobile and visible. The bad guys may not steal your laptop outright, but they don't have to; they can pop a CD/DVD into your drive — an evil little disc full of specialized tools that can load in moments and get busy while you're in the washroom or at the airport security counter. Then the blackhats retrieve and pocket the disc and go on their merry way. Without even knowing it, you're in trouble; your laptop has just been cracked.

Much the same thing can happen, even in your workplace, with desktop systems. Some criminal hackers deliberately get hired to businesses that have access to the accounts they're seeking to hack. They can bide their time, waiting for access to the targeted machines. What they want to do may take only a minute or two at most. Some criminals know to look and act the part when they make "casual" visits to your workplace. Some people like to avoid making a scene when they encounter a strange person on the premises, and a visitor who sounds and appears genuine could easily get waved by. (We've included a sidebar in this section that chronicles a couple of physical-security snafus.) Though physical security may not be a priority in a targeted company, lax procedures can be a royal road to rootkits.



The best security countermeasures against criminal hackers and thieves are primarily deterrents. If you make it more difficult for them to do their dirty work — difficult enough that they don't see you as an easy target — they may pass you by. Security products are available to assist individual users, as well as small or medium-size businesses, to limit encounters with crime. Depending on your specific needs, gearing up to meet this need doesn't have to be expensive.

The bottom line is that if you want to protect your investments effectively at home and at work you have to do it all the time. A security policy is not a formality — it's a safeguard. Create it and apply it consistently.

Using limited-access user accounts

Threats are always evolving; even tried-and-true security measures are no longer the cure-alls they used to be. For example, using limited-access user accounts (discussed in Chapter 4) instead of Administrator accounts on a

standalone system was touted online — at least until quite recently — as a way to prevent hacker attacks.

These days, a limited-access account can severely slow down an amateur, but it won't stop a sophisticated criminal hacker — nor will it stymie a rootkit. Limited-access accounts can still limit the effects of trojans, viruses, spyware, and other forms of malware because these can only load according to the parameters and permissions of the account you're using — and if it's a basic user account, that makes the pickings pretty slim for the blackhats. If you surf online using an Administrator account, all we can say is . . . *why?!* You risk the full effects of malware that way, even if you use protections. And it only takes one infestation to ruin your day.



This much of the traditional wisdom is still valid: Use limited-access accounts *with strong passwords* for your virtual travels online — and keep Administrator accounts strictly for doing configurations offline and taking care of programming needs in a secured environment. If you must be online with an Administrator account, keep it brief and surf safely.

Limiting access on networks

You can protect your network from the scourge of malware by limiting access on networks in much the same way you limit access to certain privileges — by setting up limited user accounts — except the process is more complicated when you're protecting a network. Networks have default, built-in, and special user-domain groups — all of which are different from local or standalone accounts. That means a lot of detailed configuring.



Users may belong to several different groups. It can be hard to keep track of every user and the groups to which they belong. With the `whoami.exe` tool from the Support Tools folder on the Windows XP Professional Installation CD, you can see the contents of the access token being used by a logged-on user. This glimpse can help you determine whether the user has all the correct privileges applied to the groups to which s/he belongs.

You can apply special limited-access accounts for users by creating a Restricted Groups policy using a security template. *Restricted groups* consist of multiple user accounts with limited memberships, access, and privileges to other groups. A restricted group could be set up that provided users with only limited-access accounts online — and you could make only a few memberships to other kinds of groups available on the network.

Using the Restricted Groups policy with a security template, you can ensure better protection for your network — whether the threat is from rogue users or malware. Greater control over the memberships of your multiple groups means that only the users included in the Restricted Groups policy can

have access on your network. To put this arrangement in place, log on as an Administrator and follow these steps:

1. If you haven't already done so, add the Security Templates Snap-in to your Microsoft Management Console.

Just follow the steps in the “Getting snappy with a snap-in” section earlier in this chapter, substituting the Security Templates snap-in for the Security Configuration and Analysis snap-in mentioned in that list.

2. In the left pane of the MMC, expand the Security Templates heading, then expand the template path folder, then the template you want to use (such as `securews.inf`).

3. Right-click Restricted Groups and choose Add Group.

4. Enter the name of the group for which you want to create a Restricted Groups policy, and then click OK.

The Configure Membership For dialog box appears.

5. Click Add to include members in the group.

You can also add the groups in which you want to include this restricted group as a member.

You can transfer restricted group entries from one template to another by copying and pasting within the Security Templates, MMC.



Boneheaded physical security implementations

About ten years ago, Larry was doing security work near a local hospital. One day he got lost in this rather large facility, and wound up in an entire wing of classrooms, offices, and workshops — with not another human being in sight. None of the doors were locked. He kept checking the rooms hoping to find anyone who could offer directions. The offices and classrooms had many computers — even laptops — strewn around. The workshops had expensive hand tools and other equipment. He finally did make his way out of the maze, and promptly reported his findings to the hospital's security office. Apparently they didn't listen; about three weeks later, a moving van rolled right up to the front door — and men in coveralls cleaned out the rooms in that maze, taking all the computers,

tools, and equipment. No one had bothered to check those men's credentials; they looked like they belonged. Later it was discovered that they were thieves. (Duh!) The hospital had a security policy in place — but had chosen to ignore it as “inconvenient.”

They weren't the only ones to make that mistake. In another instance, a man dressed in coveralls showed up at an office building every day, pushing a cart. He would enter, collect computers, put them on the cart, and leave the building. He was caught only after having done this for over a month. People in the building assumed he was a repairman coming to “fix” their computers. No doubt that building has better security now.

Making a business security plan

It makes sense to start with the brass tacks of physical security. Involve yourself — and your employees — in developing a security policy to safeguard company equipment. The troops on the ground can provide insights and suggestions based on their work experience. A good policy is one that includes everyone actively, rather than by imposing restrictive rules from above.

Start with a strategy — a comprehensive security plan for what you want to do with your equipment, taking into account budget and proprietary needs. Some companies require their employees and departments to take responsibility if the equipment they are issued is lost or stolen — especially if the security devices provided aren't used. Most thefts of desktop computers occur with common burglaries, but laptops are most often stolen as crimes of opportunity. Laptops are becoming increasingly popular as items of theft — after all, they're easy to move, conceal, and dispose of — but that evil little cracking disc mentioned earlier in the chapter can be just as damaging a means of theft.

The best strategies provide a thorough plan for protecting equipment, including these measures:

- ✓ **Equipment classification, cataloging, and labeling.** You need an inventory of your equipment — including clear definitions of what purposes and capacities it serves. Try these questions for openers: What items would be expensive and catastrophic if lost or stolen, which ones are most important to company goals, which ones are in public access areas? Labeling your equipment means putting anti-theft tags on each piece to promote recovery and to make it hard for the bad guys to resell the stuff. Secure-It offers an online source of examples you can use to develop asset-tagging systems:

www.secure-it.com/shop/product_info.php/cPath/42/products_id/34

- ✓ **Physical security devices to protect the equipment whether in or out of the workplace.** You can physically lock down your workstations so they can't be tampered with or stolen; special enclosing and affixing devices provide a strong deterrent to thieves. Metal cases with locking cables or chains are best for securing laptops to any work area. These devices can also have built-in alarms to prevent tampering — or trigger an alert if somebody tries it. Cabinets can provide secure storage for laptops not in use. Here are some online resources:

- **Securtech:** www.securtech.com/cgi-bin/STORE/store.cgi
- **Safemark:** www.safemark.co.uk

- ✔ **Software for data protection, access control, and recovery.** What if the worst happens? Is the data and information on your computers secure from criminals? Passwords can be cracked, providing any type of remote access to your documents or network. What can you do about it? Software solutions exist that can block the illegal use of your computers and networks. These links offer more information:
 - **CyberAngel:** www.thecyberangel.com/sohoapp.html
 - **CompuTrace:** www.absolute.com
- ✔ **Employee and personal training in security awareness.** Awareness — especially security awareness — is the key to fending off blackhat attacks. Employees sit behind — and sometimes in front of — your firewall and other security applications. Properly trained, they can prove to be your most powerful asset in securing your computing business. If you want your employees to understand and properly implement a security strategy (or want to brush up on that yourself), design a security-training program that matches your business model, network size, and user requirements. This does not have to be time-consuming; a half-day seminar could suffice. Here the everyday security concerns can be highly effective to hash out and implement — especially in relation to e-mail usage, passwords, access control, Internet downloads, social engineering (those innocent-looking guys in the coveralls), phishing (search Google for the term phishing if you're unfamiliar with it), home and remote Internet use, removable media, controlling software updates, and making sure everybody knows how to use the anti-theft devices and techniques you provide.

Every one of these concerns you nail down can help deflect rootkits. Fortunately, it's also possible to outsmart them — and that's next.

Fooling Rootkits with Virtual Operating Systems

Even these days, rootkits are pretty simpleminded; they can't discern the difference between a *virtual machine* (a part of the operating system set up to behave like a standalone device) and a real computer operating system. A virtual machine's "hard drive" is really just a window; if the VM becomes infected with any kind of malware or a rootkit, it can be reset to pristine and pure condition in a few minutes. Rootkits roll off VMs like water off a duck's back. No wonder virtual machine (VM) software using Windows operating systems could become a wave of the future.

Rootkits are easily repelled by VMs in most instances. Two major VMs are

- ✓ **Microsoft's Virtual PC:** Microsoft's Virtual PC works best with Microsoft operating systems — but not at all with Linux variants. System requirements: 400 MHz Pentium II or better, 32MB of RAM, Windows 2000 or XP.

www.microsoft.com/windows/virtualpc/default.aspx

- ✓ **VMware Workstation:** VMware's Workstation works well with Linux variants. System requirements: 500 MHz Pentium II or better, 128MB of RAM, and Windows NT 4.0, 2000, XP, 2003, or various Linux distributions.

www.vmware.com/products/ws/

Both products have pros and cons — but they work well as virtual machines with few bugs. Both run on Windows 2000 or XP.



Be sure your machine has plenty of RAM. The VM runs on the unused portions of your RAM and CPU cycles. At least double the RAM you normally use should suffice. Using at least a Pentium 4 processor seems to yield a faster response as well.



Virtual machines can even be used as network servers. Both Microsoft and VMware provide server versions of their software. Large networks (in practical terms, those with over 500 computers) are using VM to help handle their loads.

For more information on virtual machines, please refer to

http://en.wikipedia.org/wiki/Virtual_machine

Planning Your Defense Against Rootkits

Rootkits are evolving and developing rapidly. The ones that were so familiar a year ago are fast becoming obsolete as new forms, such as those using DKOM, replace the ones that used hooking techniques. No, we're not indulging in rootkit nostalgia here. The old rootkits are still around — and they're still dangerous — but the newer ones are even more insidious and difficult to detect.



Okay, the newer rootkits can evade file-integrity checks — but we still advise you to keep doing those (see Chapter 9 to learn about file-integrity checks). Regular file-integrity checks can help you identify clues to the malware or weaknesses that would otherwise give blackhat hackers a way to gain access to your computer or network.

A *baseline* is a record of what's considered normal use for each particular computer — with a system snapshot and hashes for all essential system and network files. Another step in the direction of a happier future is to prepare your own recovery discs when your system is clean and running properly.

Establishing a baseline

Establishing a baseline is about using file-integrity applications to monitor essential system and network files for any abnormal behavior or unexpected changes. You create your baseline when you run the program on a computer that you know to be clean of malware and running normally. The program makes unique identification keys, or fingerprints, with complex cryptographic hashes for each essential file and folder. The keys are created by running the files and folders through an algorithmic filter such as SHA-1, MD5, SHA-2, and so on. These can then be compared to standard or established hashes for those files as well as to your baseline. Later scans are then compared to your baseline information to check for unknown or untoward changes. When any of these files and folders are changed along expected lines, they can be updated into the baseline to keep it current. (See Chapter 9 for more information on hashing and verifying file integrity.)

The idea is to keep a running tab on all essential system and network files to see if they have been changed by an exploit or malware. Instead of directly detecting the malware itself, you're checking to see whether it's left footprints (as discussed in Chapter 8). This process can come in handy when you're dealing with new forms of malware. Successive scans are run, comparing the current files and folders with the baseline, revealing any and all changes to them. Although these can yield false positives from time to time, they're effective at providing information on possible infections and hacking attempts.

Corrupted files and folders can then be replaced with backups comparable to those used to establish the baseline. Files or folders that have been added — or that don't belong — are then easier to spot and can be removed. You need to establish this baseline to have something to compare against later, such as when you're trying to sniff out signs of rootkits or malware as discussed in Chapter 8.



Similar to regular backup procedures, file- and system-integrity scans should be done with the computer or server offline. This prevents files from being added or changed by online influences. For example, if you have a trojan on your computer and do a file-integrity scan while online, you provide the intruder with the opportunity to connect with its creators — and the first thing it'll want to do is obfuscate your efforts.

Before blackhat hackers can put rootkits on your system, they need to make changes to your system files to enable backdoor access. If you run regular file and system integrity scans, chances are you'll catch those changes before they disappear from view (which would happen after the arrival of a rootkit).

Preparing Recovery Discs

In Chapter 3 we provided some information on using the Windows Backup utility. That utility is fine, provided you want to back up to a tape drive. However, most of us have CD/DVD ROMs on our computers. Floppy drives — even CD-ROM drives in some cases — are fast becoming obsolete; they just can't handle the sheer volume of software measured in gigabytes. You can use Windows Backup with other kinds of USB plug-ins or Flash drives, but be sure to check their capacities before using them; Windows Backup doesn't support using more than one copy each of these drives. This happened to Larry when he used his Windows Backup utility for the first time many years ago. He tried to do a backup of his hard drive with CDs. It worked fine on the first CD but just wouldn't do any more. After much searching in the Windows Help files he discovered that the backup procedure did not support multiple copies in sequence. He wondered about the point of such an inadequate backup utility, and went looking for alternatives. He found them in Karenware's Replicator and the Acronis True Image.



You could use these two applications to back up your system: for smaller jobs (such as for the My Documents and Settings folder), use The Replicator; for really big jobs — such as your entire hard drive — use the Acronis True Image 9.0 Home. Both of these programs are available on the *Rootkits For Dummies* DART-CD.

You may have received *recovery discs* with your computer when you purchased it — bootable discs designed to restore your system after a crash. In normal difficulties and problems, these can be useful — but after infection by a rootkit, they quickly become useless because you cannot use them to recover from a rootkit infection. Rootkits may cripple the system — but that's not same as a system crash. A rootkit could change the model number of the computer (for example) thus rendering the recovery discs useless. Most network servers make regular backups of both the server and selected workstations on a regular basis, especially to separate media. They use these in case of emergencies, to reduce their downtime and to restore both systems and the network. Single users, faced with the threat of rootkits and their payloads, will need to create their own recovery or restoration disks — just like the networks do.

Simply backing up your operating-system files and software may not be enough anymore. Some rootkits leave systems in a completely unusable state, even though their active menace may have been removed. Having a full ISO image of the system before it was infected makes later restoration much easier.

The Replicator is exceedingly easy to use — and will copy files to any form of media, or to any other computer on a network, provided you have the equipment. Windows XP won't allow you to create a copy of a large folder such as My Documents in the same archive, but you can do it with the Replicator.



The Replicator uses a time and date selection in the interface; if you want to set the time for a single job now, enter today's date and a time in advance.

Depending on the size of the job, you may need to backup to a selected folder on your hard drive with the Replicator first — and then copy it to CD/DVD with your burner. The Replicator can do a backup of your entire hard, if you wish. You can have it do regular backups of selected files or folders, on a timely basis — automatically — backing up only what has changed since the last job. More detailed instructions are provided in the Appendix and in Bonus Chapter 2.

If you want make an ISO image backup of your entire hard drive with Acronis, you'll find detailed instructions in the program's manual (you can get copies of their manuals at the address below). To make this backup on separate media, you'll need a DVD (*not* a CD) or a UDF burner. (CDs just don't hold enough data to do the job.)

www.acronis.com/homecomputing/download/docs/

The purpose for having full backups of your system and drives from an uninfected time allows you to almost completely restore your system after you follow the instructions provided in Chapter 11. The alternative is a fresh installation of your operating system — which entails losing all your files and work because your hard drive has to be totally erased to be completely cleansed of *all* rootkit traces. Here's where a little pre-emptive paranoia can be useful; prepare for disaster *before* it happens to avoid such potential headaches down the road.

Part III

Giving Rootkits the Recognition They Deserve

The 5th Wave

By Rich Tennant



"A centralized security management system sounds fine, but then what would we do with all the dogs?"

In this part . . .

By examining rootkits, you will learn more about how the enemy, *blackhat hackers*, use rootkits to further their goals. Sun Tzu said, “Wage war with surprise moves.” A rootkit is an especially rude surprise, and a most stealthy one at that. It’s a weapon and a secret agent rolled into one. You need to know how this weapon works, what it does, and how it persists before you can defuse and remove it. A wise soldier prepares ahead to meet the coming foe.

In this part, you get to discover why rootkits are the most dangerous software ever devised. Know thine enemy!

Chapter 7

Getting Windows to Lie to You: Discovering How Rootkits Hide

In This Chapter

- ▶ Understanding privileges from a rootkit's point of view
 - ▶ Getting familiar with rootkit hooking tactics
 - ▶ Checking out alternative, more secretive rootkit technologies
 - ▶ Looking into the future of rootkit shenanigans
-

The presence of a rootkit on a computer or network implies a violation of the very framework of the system you have come to know and trust — and that's a security breach of the highest magnitude. No wonder rootkits have to hide to be able to stick around long enough to perform the nefarious functions they were designed for. (There's more about that in Chapter 1.)

How can a rootkit be so all-powerful that it actually gets Windows to lie to you? Well, that power comes from taking advantage of the flexibility and versatility that were built into the Windows architecture — a rootkit just manipulates those features to suit its own needs. That's why this chapter investigates how rootkits trick the operating system into becoming an unwitting accomplice to their deception.

Discovering How Rootkits Hide and Survive

In effect, rootkits are invisible to most traditional malware scanners. They hide by getting the operating system to lie for them or (in the case of scanners) to report falsified results. They intercept and filter system information so the operating system's output excludes any indication of the rootkit.

Whenever a user or program queries the system, the normal assumption is that the information the operating system returns is valid. A rootkit can violate this basic trust — and when that happens, the entire premise of accuracy and reliability that we have come to depend upon gets turned completely upside down. A well-crafted rootkit makes sure the operating system returns only doctored information that doesn't betray its existence. A rootkit hides itself — and all the infected malware components associated with it. It's able to accomplish this by obtaining Administrative access to a system. Once a rootkit gains control of the Administrator account — the Big Kahuna of the Windows operating system — it can then control what the system does (and what it reports).

When a rootkit gets hold of administrative privileges, it can secretly go about changing and hiding vital system components such as files, directories, ports, the Windows Registry, and even the code of the operating system itself. One critical way the rootkit maintains concealment is by *hooking* system function calls (more about those in a minute) — altering the data returned by system utilities and scanners so no evidence of system tampering is revealed.

Some rootkits have what amounts to a self-preservation instinct: The malware associated with such a rootkit (its *payload*) employs a feedback mechanism to ensure the rootkit's survival. When a rootkit command is executed on a compromised system, a component of the malware payload examines other system commands to ensure that the rootkit is fully intact. If it is not, then the rootkit reinfects the system until it's in place and functioning again. In a similar manner, the payload may establish a backdoor and monitor traffic on that port. If the traffic flow doesn't match that of an infected machine, the remote attacker reinfects the host computer (server).

When hackers claim their prize, they sometimes go to extraordinary lengths to stake their claim. Though not associated with a rootkit — yet — one recently discovered threat actually comes with Windows security patches in tow. That's right — once it compromises the system, it plugs the very vulnerability that allowed it to enter so that it can have sole ownership of the computer. Another recently discovered threat that is not known to be rootkit related — again, yet — installs a counterfeit copy of a highly regarded antivirus on the system it has compromised. Once it's on board, it cleans up its new host to prevent other would-be intruders from gaining entry.

Chilling? Creepy? You bet. A hacker will stop at no extreme to own a system, and a rootkit is the ideal tool to allow them to maintain complete and unfettered access. All rootkit modifications — whether at the level of the individual user or the OS kernel — have *one* aim in mind: Conceal the presence of the rootkit and its payload on the infected computer so the payload can proceed with its dirty work without interruption.



Rootkits exist not only on Windows machines but also on other platforms, such as Mac OS, Linux, and UNIX. Because Windows is by far the most popular operating system, most malware writers have chosen to concentrate on Windows so they can get the most bang for their buck.

The remainder of this chapter is a rogue's gallery of rootkit functions — the ways they get access to your system, the different types of rootkits, and just how they hide.

Keys to the Kingdom: Privileges

The concept of *computer privileges* is the key to how rootkits operate — and to how they're classified. Start with a practical fact: Not everybody uses a computer in the same way — especially on a network. So network administrators assign different levels of privilege to different users; the idea is to provide appropriate capabilities (and only the needed ones) for doing all those different jobs. The Administrator account gets to say who can do what.

The most efficient way to imagine computer privileges is as a series of concentric rings around a central point — like a simplified representation of the solar system. Here the “sun” is the kernel of the operating system; and the planetary “orbits” are the rings that represent the different privilege levels of the operating system. The ring closest to the kernel (analogous to the closest orbit around the sun) is referred to as *Ring 0*. Any component that has Ring 0 privileges is operating at kernel level — the highest level of privilege. As you follow the rings or circles farther out, their distance from the center increases — and the farther out they are, the fewer privileges they have.



In reality, Windows systems have fewer rings than the solar system has planets (even if you don't count Pluto). In case you're wondering, Ring 1 and 2 privileges are not used in current Windows platforms; Microsoft used them last in the 1990s, in an operating system called OS2.

If an application wants to access a system resource such as the hard drive or memory, it must communicate with the operating system to do so. How direct or indirect this level of communication is, is determined by the level of privilege a given application has. Rootkits with higher privileges have an easier time compromising your system.



Rootkits need administrative access to infect a computer, so two very basic ways to limit your risk of infection are (1) to log on as a limited user and (2) use strong passwords to protect all your user accounts.

An example of a rootkit snafu

The University of Connecticut (UConn) computer system housed a rootkit for nearly two years before it was noticed, according to a June 27, 2005 article in *eweek*. The sheer scale of possible trouble such a rootkit could have caused was considerable; it was detected on a server that contained identity information for 72,000 people — students, staff, and faculty.

Luckily — *very* luckily — in this particular case, there was no leakage of sensitive information

because the rootkit's attempt to install a backdoor failed. Still, the incident serves as a warning to network administrators and IT professionals, and stresses the importance of having adequate recovery and security measures in place. Although the UConn rootkit made the news (basically because they reported it), no doubt there are other servers that have fallen prey to similar attacks but have managed (so far) to escape the public spotlight.

Knowing the Types of Rootkits

Rootkits are classified in two basic ways — as *user-mode* or *kernel-mode* — depending on the scope of the effect and whether they exist only in memory or have written changes to disk that enable them to survive a reboot. Rootkits can reside at the level of user accounts (Ring 3) or at kernel level (Ring 0) — and they can be either persistent (able to survive a reboot) or non-persistent (that is, they disappear at reboot). Table 7-1 provides a basic summary of rootkit types.

<i>Type</i>	<i>Privilege Level (Ring)</i>	<i>Scope of Action</i>	<i>Persistent or Non-persistent</i>
User-mode rootkit	Low (Ring 3)	Localized effect	Can be either
Kernel-mode rootkit	High (Ring 0)	System-wide (global) effect	Can be either

The following sections give you a closer look at user-mode versus kernel-mode rootkits, and at persistent versus non-persistent rootkits.

User-mode versus kernel-mode rootkits

Rootkits are classified by the mechanisms they employ to infect a computer — and by the *scope* of their action (that is, the extent of dirty work they can do). Scope is, in turn, determined by the privilege level a rootkit can obtain for itself. Higher privilege levels mean a greater scope.

Whatever their mode, here's what every rootkit knows: All programs communicate with each other through function calls. Depending on a program's privilege level, it can operate either by making function calls directly to the kernel (if it's at Ring 0) or from a more restricted level of privilege (at Ring 3, the level accorded to user programs). Privilege level determines whether a rootkit operates in user mode or kernel mode (more about that in the next section). Presumably all rootkits would love to reside at Ring 0 — but they'll take whatever privilege level they can get (which really translates into whatever their author is able to code) to operate effectively.

User-level rootkits

User-level programs, orbiting out there at the lower levels of privilege, must use the *application programming interface* (API) to make requests for operating-system resources. These *system calls* go to the kernel indirectly, through user-level dynamic link libraries (DLLs). The DLLs translate user-level API calls into calls the kernel can understand. In effect, user processes must operate through a middleman to talk to the kernel.



Now, a user-level program isn't self-sufficient. Some of its basic needs are provided by the operating-system kernel instead of the program itself — for example, such functions as reading or writing to disk, displaying a window, or printing a document. A user program must make *system calls* to the kernel to request that it perform those actions. For example, if a user clicks an OK button, Windows translates that input into a system call — and then asks the system to act upon it. The language that Windows provides to accomplish this communication between kernel and user program is called the *application program interface* (API). To operate more efficiently, each user program builds its own its unique table — which eventually contains the addresses of all APIs or system functions that it needs the kernel to perform to complete its execution. This table is called the *Import Address Table* (IAT), which makes another (more detailed) appearance later in the chapter. The IAT is part and parcel of a program's executable (EXE) file (or image) that is loaded into memory.

The indirect calls that user-level programs must use limit the effect of user-mode rootkits (which can run only within the confines of another application, or as separate user-level programs). Unable to interact with the kernel directly, limited to what a user program can do, user-mode rootkits exert only a localized effect. They are kept at arm's length away from the kernel, in

user program address space, and must make requests for kernel services through system calls (see the “Hooking to Hide” section later in this chapter), user-mode rootkits can be detected by some conventional security programs, (and nearly all rootkit detectors) that run in kernel mode. Despite that disadvantage, user-mode rootkits are easier to program than their kernel-mode cousins — and are therefore less likely to crash or hang the operating system, so they can operate as long as a kernel-mode rootkit detector isn’t looking for them. The trade-off (from the bad guys’ point of view) is that a user-mode rootkit is less powerful, and a lot of work has to be done to achieve the desired global effect. To be effective, a user-mode rootkit must find a way to alter the tables of *every* executing user program.

Kernel-level rootkits

The operating system *kernel* is the software equivalent of the central processing unit — the “brain” of the operating system, and its most basic component. The kernel provides the buck-stops-here level of control and functionality for all programs that run on a computer. It maintains and manages many vital system resources and functions — such as memory, security information, and process scheduling — and facilitates communication between software and hardware. The kernel therefore has a global scope — access to the entire operating system, hardware and all applications. All programs must interact with the kernel in some way — and if a rootkit gets control there, it’s in the driver’s seat.

Unlike user-mode rootkits that operate at Ring 3, kernel-mode rootkits that operate at Ring 0 interact with the kernel more directly — by intercepting *native* (or *kernel-level*) APIs. The kernel has global access to every nook and cranny of the operating system — so inside the kernel is a comfortable place for a rootkit to be. From there, it can access any memory location and any hardware (worse) it can substitute its own code for that of the kernel’s or modify the critical data structures that the kernel relies on to keep track of its activities. Kernel-level rootkits can, therefore, exert a global impact on the entire system — which makes them potentially a lot more dangerous and insidious.

Kernel-mode rootkits install a device driver to obtain access to kernel-level privileges. Once ensconced, the rootkit driver redirects system function calls so its own code is executed instead of kernel code. (For more about how this process works, see the “Hooking to Hide” section later in this chapter.)



Luckily, it’s inherently difficult to implement a kernel-mode rootkit without upsetting the delicate balance of the system kernel. For that reason, kernel-mode rootkits often reveal their presence by causing system instabilities or computer crashes (especially if the rootkit has been sloppily programmed).

Persistent versus non-persistent rootkits

One other essential distinction between rootkits is whether they can survive a reboot. If they can, they're called *persistent*; if they can't, they're *non-persistent*. A closer look at each of these types is the next order of business.

Persistent rootkits

In order to survive a reboot and attain *persistent* status, a rootkit must physically alter the contents of the hard drive. A persistent rootkit does this bit of nasty magic by residing on the disk and adding an autostart entry to the Registry — that way it's loaded into memory and executed automatically every time the computer is started. Although these physical changes to the disk do present opportunities for detection, unfortunately many rootkits are still undetectable to traditional malware scanners and system utilities. What escapes attempted analysis is a range of hidden clues (more about those in a minute) that would alert most dedicated rootkit scanners to a rootkit's presence.

Non-persistent rootkits

Non-persistent rootkits present even fewer detection opportunities to scanners because they exist in memory only — and disappear if there's a reboot. Programs that merely scan physical storage media have no chance of detecting a non-persistent rootkit — after all, it leaves no part of itself behind. True, non-persistent rootkits may seem less threatening to those of us who reboot regularly; restarting the computer should effectively eliminate the threat. Well, it does — but there's a catch: If the infected computer is a network server — connected to hundreds of client computers — rebooting is typically a much rarer event. Memory-resident rootkits take advantage of this fact, and stick around for as long as they can. Consequently a network may harbor an undetected infection for an extended period of time. The very fact that non-persistent rootkits leave no physical clues makes them a lot harder to detect.

Hooking to Hide

One technique a rootkit employs to alter the normal execution path of the operating system is known as *hooking* — intercepting system function calls and adjusting their results to deny the presence of the rootkit. Hooking diverts normal program flow, to rootkit-supplied functions instead of legitimate system functions. The Windows operating system was designed to be

very adaptable and flexible — which probably seemed like a good idea at the time, but as such provides many possibilities for rootkits to “hook up” to system resources.

In the following sections we’ll discuss how rootkits hide by attempting to make themselves invisible to the target operating system — and we’ll give you the skinny on the various types of hooking.

How hooking works

Both user-mode and kernel-mode rootkits employ hooking techniques to filter the results returned by the operating system and camouflage their existence. In effect, they get Windows to lie for them and perpetuate the illusion that you have a clean system. Rootkit hidden files Registry entries, processes, and ports — will be invisible to most system-analysis and scanning programs. That’s because such programs rely on the data provided to them by the operating system to produce their results. For example, if you use (say) the Task Manager to display a list of active processes, the rootkit processes will be excluded from what you see. The same will happen if you try to locate rootkit files and folders using Windows Explorer. Likewise, Regedit won’t detect any of the rootkit-installed Registry keys; Netstat won’t see any rootkit ports. It’s like being in an alternate universe where none of your trusty tools can actually be trusted; Windows blindly becomes a partner to a rootkit’s subversion techniques.

The Windows operating system uses many data structures (tables) to store and keep track of critical system information. These tables can be hooked, replaced, and generally corrupted by a rootkit. User-mode and kernel-mode rootkits both use hooking, but what they’re permitted to hook — and the tricks they use — are defined by the privileges accorded to them:

- ✔ Some rootkits sport a type of built-in enemy radar. They temporarily drop their hooks when they detect a rootkit-scanning program poking around in their vicinity. Then no hooking abnormalities show up in the scan results — and the rootkit remains undetected. They may possess an uncanny intelligence — one of the newer rootkits completely disables all tools capable of detecting it, and even blocks Web access to security sites that provide removal tools.
- ✔ User-mode rootkits can only hook data structures (tables) in a user program’s address space (the IAT and EAT), so the scope of their effect is limited. They may also insert jump instructions into user-level APIs; the result is to redirect system calls to the rootkit’s replacement functions.

- ✔ Kernel-mode rootkits hook tables (data structures) used by the kernel — such as the SSST and the IDT — so a kernel-mode rootkit's effect is system-wide or global. (These tables and the various types of hooking are described in the next sections.)

Knowing the types of hooks

Two types of hooks exist, including privileged hooks and unprivileged hooks. User-level rootkits use unprivileged hooks exclusively while kernel-level rootkits use the more efficient privileged hooks. Here's a list that includes the various types of hooking associated with each:

- ✔ **Unprivileged hooks:** User-mode rootkits run within the confines of user program address space and use *unprivileged hooks* to redirect program flow. The hooks are called unprivileged because they do not have Ring 0 privileges (they have Ring 3 privileges only); the rootkit exerts its influence by operating within the memory-address space that has been allocated to another program.
- ✔ **API hooking:** User-mode rootkits intercept the API (Application Program Interface) calls that user programs use to communicate with the kernel using unprivileged hooks. They hook (modify the addresses of) APIs in the Import address table (IAT) of user processes, so they point to rootkit functions instead of the Windows API functions.

Most rootkits (both user- and kernel-mode) use API hooking to make sure the operating system returns only filtered results, which omit any indication of the rootkit or its payload.

User-mode rootkits can only modify tables that belong to user programs. Each user program has its own unique IAT that it assembles to indicate the functions it needs the kernel to perform; to be effective, a user-mode rootkit must find a way to alter the tables of every executing user program — and also monitor for any new programs that start up so it can hook the relevant APIs used by those programs as well. A rootkit doesn't need to hook every API, only those that may expose it. For example, it may hook the APIs that Task Manager uses to display its list of active processes, or the APIs that are used by Windows Explorer to display files and folders. Because it has to hook these same APIs for every active process, user-mode rootkits are much less efficient than kernel-mode rootkits. Kernel-mode rootkits can achieve the same effect by hooking a single structure that all user programs use.

- ✔ **DLL injection:** User-mode rootkits use a method called DLL (Dynamic Link Library) injection to implement unprivileged hooks. To get a better handle on what DLL injection is, and how it's used in unprivileged hooking, check out the next section to get a closer look at what a DLL is.



DLLs and the rootkits that love them

A DLL is a type of small program — often it is a building block for other programs — called an *object module*. Some DLLs are application-specific and reside in the application's program folder; other DLLs are provided by Windows itself — and can be shared by all programs. Windows maintains its own collection of small shared program modules used to perform identical system functions for all applications. These interchangeable DLLs reside in the Windows system folder.

A larger program calls a Windows DLL program when it needs to perform the function that the DLL provides. For example, the DLL may provide a print function for a word-processing program, or allow you to write to (or read from) your hard drive. Windows DLL files are not part of the program's executable file; instead, they're linked in at runtime and loaded into RAM as needed. Utilizing shared DLLs this way conserves system resources; executable files can take up less space because they can “borrow” some of their basic functionality from DLLs that aren't loaded into memory until they're needed. DLLs often take the form of drivers — and drivers are a favorite tool of rootkits (as we shall see later in the chapter).

If an application has to use a particular version of a DLL, the application installs that version of the DLL in its own program folder to avoid conflict with the version provided in the Windows system folder. The application then decides which functions it must import from the DLLs that are available to it at runtime.

Static and dynamic DLLs as rootkit targets

Static DLLs, as opposed to dynamic DLLs, are application-specific; because they aren't used by other programs, they can be linked or compiled directly into the application's executable file. Because they are contained within a program's executable file when it is loaded into memory, unlike dynamic DLLs static DLLs are not vulnerable to rootkit modification — unless the rootkit replaces the entire DLL file on disk — but that method is easily detectable and rarely used nowadays.

Windows uses a system called Windows File Protection (WFP) to ensure that the DLL files in its system folder are not overwritten by different versions supplied by application-program DLLs. This means you cannot replace the system version with the user-program version, because Windows will immediately overwrite the new third-party DLL with a bona fide version that it maintains in a system backup folder.

The prize DLLs at kernel and user levels

The `Kernel32.dll` and `User32.dll` files are crucial; they supply system functionality to user programs. For example, `Kernel32.dll` contains APIs that manage memory, handle program interrupts, and take care of input/output tasks. (Lots of potential mischief to be done there.)

You may become aware of the `Kernel32.dll` operation when an `invalid page fault error` or `kernel error` message is generated. The Windows dynamic link libraries are loaded into a reserved area of memory at system startup. Because user programs are not allowed to access the kernel memory space directly; they have to use API calls when they need to perform a function that only the operating system can provide. This difference in access maintains the concept of privilege — user programs have no direct contact with the kernel.

`Kernel32.dll`, `User32.dll`, `Gdi32.dll`, `Ws2_32.dll`, `Advapi32.dll`, `Wininet.dll`, `Rasapi32.dll`, `Urlmon.dll`, and `Netapi32.dll` all contain user-level APIs that can be called to execute specific functions for user-level programs. Because the APIs must be retrieved from user-level DLLs that are not part of the application, these DLLs and their APIs are considered to be “imported.” The functions within these imported user DLLs will in turn call the kernel or native APIs required to process a given user-program request. `Ntdll.dll` is a bit of an oddball; it provides a native API interface directly to user-mode programs, even though `Ntdll.dll` still resides in the portion of memory reserved for user programs. User-mode calls to `Ntdll.dll` APIs cause `Ntdll.dll` to execute a `SYSENTER` (analogous to generating an interrupt in older versions of Windows) which immediately transfers control to the kernel. It is in the kernel that the APIs are processed by executing the appropriate system services — so `Ntdll.dll` basically functions as a user program’s gateway to the kernel.

An application maintains a pointer (which is just an address) to each imported user-level API it needs within a special data structure called the Import Address Table (IAT). The IAT enables an application program to locate and execute a system function whenever it needs it, and by handling it this way, Windows makes its operations more streamlined and efficient. The IAT may contain entries that point to any of the DLLs we’ve listed here. An application assembles its IAT before runtime when it knows the name of the APIs it needs — but it doesn’t yet know their addresses. What it does provide are the names of the APIs that it needs, along with the names of the parent DLLs that supply them. Windows uses that information to load the required DLLs, and then it is Windows that computes and fills in the program’s IAT with the correct API addresses.



More on DLLs

In reality (say, program executables), each user DLL has an IAT for any APIs it may be importing from other DLLs such as `ntdll.dll`. Windows checks the IAT of each user-level DLL to see if it's importing any functions from any other DLL. If

it is, then Windows will load that DLL and overwrite the IAT of the calling DLL with the actual addresses of the APIs in the called DLL, using the EAT of the called DLL. If called DLL is `ntdll.dll`, it will already be loaded into memory.

A rootkit is able to fool Windows by overwriting the address information in a program's IAT with address information that points to its own code, thus mimicking a variation of this last step that Windows has already completed. It accomplishes this feat in the following way: A user-mode rootkit substitutes its own APIs for the normally imported ones by injecting its own DLL into the program's address space. The rootkit can then selectively hook the IAT by replacing legitimate pointers in the IAT with pointers to its own functions. When a user program calls the one of these hooked APIs, execution is redirected to the rootkit's replacement function (located in the rootkit's *injected DLL*, so named because the rootkit injects it into a user program's address space). The rootkit need not hook every API, but only those APIs that are capable of exposing it.

In actual practice, there is another table called the Export Address Table (EAT) that presents another user-mode API hooking opportunity. Unlike the IAT, which is located within a user program's executable file, the EAT is located within the imported DLL itself. An imported DLL's EAT stores pointers to the functions that are exported to user applications already loaded into memory. These entries correspond to those in the IAT of the user program, except for one difference — they contain actual address information for the APIs the application program needs to import. The EAT information in the DLL is used by Windows to fill in the user program's IAT with real addresses (as opposed to just names). The EAT is vulnerable to rootkit hooking, and if the EAT is hooked, the same address information will be duplicated in the user program's IAT because that's what the IAT information is derived from.

To help you pull all this information together we'll first summarize the sequence of actions that take place during *normal* program execution:

- 1. When an application program is executed, Windows loads it into memory.**

2. Windows checks to see if the IAT of the program is calling any APIs.
3. If it is, Windows loads the DLLs specified in the program's IAT.
4. Windows then computes the real API addresses using the information stored in the EAT of each imported DLL.
5. Windows overwrites the program's IAT with the actual API addresses.

Now we'll look at a practical example (using the `FindNextFile` API) that shows how a rootkit might use EAT hooking to make sure its files stay hidden:

1. First, the rootkit scans memory for DLLs that export the `FindNextFile` API.
2. Because `FindNextFile` is in `Kernel32.dll`, this really means it scans memory for all instances of `Kernel32.dll`.
3. The in-memory `Kernel32.dll`'s EAT will contain the address of `FindNextFile` and the rootkit will overwrite that address with the address of a replacement function located inside the rootkit's injected DLL (which is already loaded by the rootkit).

The influence of a user-mode rootkit's *unprivileged* hooking mechanism is localized. That's because a user-mode rootkit can only modify the data structures available to user programs — and they must inject code into each new process that starts in order to hook any APIs that might reveal its presence. Even though all user-level APIs eventually get converted to native (kernel) APIs, the rootkit can only achieve its goals using level of indirection — and that restricts the scope of unprivileged hooking.

Kernel-mode rootkits, however, overcome this restriction by using *privileged* hooks (more about those a little later in the chapter) to directly modify the data structures that the kernel accesses. The reason is simple: Kernel data structures are used system-wide — so privileged hooking allows kernel-mode rootkits to exert a global effect.

Inline Hooking: A more devious variation on a theme

Beside IAT and EAT hooking, there is another form of hooking used by both user-mode rootkits and kernel-mode rootkits: inline function hooking. User-mode and kernel-mode inline function hooking differ only in regard to what is hooked — user-mode rootkits inline hook APIs imported from user mode DLLs, while kernel-mode rootkits inline hook the native APIs functions that reside in kernel space (see the “Kernel Inline Hooking” section later in this chapter).



An example of DLL injection at work

Here's where we illustrate how a rootkit might use DLL injection and API hooks so it can fake out both the operating system and the user.

Suppose a user invokes a search program to see whether a particular file or folder is present on the local hard drive. The search program uses a couple of Kernel32 APIs — `FindFirstFile` and `FindNextFile` — to conduct its search. Now, suppose the search program is rummaging around in a folder that contains a rootkit file. The rootkit wishes to remain hidden, so it injects a DLL into the address space of the file-search program, and replaces (patches) the `kernel32.dll` addresses with addresses contained within its own injected DLL. That

way, when the `FindFirstFile` and `FindNextFile` APIs are called, the execution is diverted to the rootkit-supplied functions in the injected DLL instead of the legitimate `kernel32.dll` functions.

The malicious code lurking in the rootkit DLL for these APIs ensures that the rootkit file is not reported and remains hidden. One of the methods a rootkit can use to accomplish that is by having the replacement function simply call the `FindFirstFile` and `FindNextFile` APIs once again, and then pass control back to `kernel32.dll`. This makes the API skip the rootkit file and jump to the next file, so any search results returned will exclude the rootkit.

Inline user hooks overwrite the actual API function code implemented by imported DLLs (not just a pointer), leaving the IAT and EAT unaltered. Inside the hooked DLL, the first line (usually) of the actual code of an API is overwritten with a jump (JMP) instruction. The JMP jumps or redirects execution to a replacement function inside the rootkit DLL.



Some sneaky rootkit programmers have attempted to insert the JMP instruction further down in the API code, thereby defeating rootkit detectors that only examine the API's first instruction to establish whether an inline hook is present. Inline hooking is much more difficult to detect than API hooking.



A program called API HookCheck by SIG² can detect all types of user mode API hooking including inline function hooking. API HookCheck compiles a list of any API addresses from the IAT and EAT that lie outside the memory space of imported DLLs. The program then generates a log composed of these suspicious entries. ApiHookCheck looks for IAT and EAT patching, as well as inline function hooking. You can find more information (and a sample log) at this API Hook Check download link:

www.security.org.sg/code/apihookcheck.html

Thread injection

User-mode rootkits can also use a method called *thread injection* to gain entry into a user program's address space. Simply put, a program consists of a series of steps that are executed in sequential order to complete some pre-determined goal. A program step may call for a process to be executed. Each individual process is composed of smaller functional units called *threads*. A rootkit can insert its own stowaway thread into a user process to ensure that its code gets executed along with that of the user program. The API used to do this is called `CreateRemoteThread`. This injection method requires careful monitoring of starting and stopped processes to properly hide the rootkit components.

AppInit_DLLs injection

Malware writers can use a relatively easy technique to inject code into nearly every user-mode process: *AppInit_DLLs injection* — thereby coming pretty close to achieving the scope of a kernel mode rootkit. Most user processes link to `User32.dll`, and when they do, any DLL(s) specified by the `AppInit_DLLs` Registry value are also loaded. The `AppInit_DLLs` is defined by the following Registry key:

```
HKLM\SOFTWARE\Microsoft\WindowsNT\CurrentVersion\Windows\
AppInit_DLLs
```

The `AppInit_DLLs` provides a very convenient method of introducing a malicious DLL that can affect almost every user process. All the bad guys' code has to do is change one Registry key value and create an infective DLL that is then injected into every running process that normally links to `User32.dll`.



A trojan that uses `AppInit_DLL` injection

Though not a bona fide rootkit, A Cool Web Search (CWS) variant was one of the first infections to utilize `AppInit_DLL`'s injection in early 2003. The particular CWS variant involved was dubbed "the About:Blank" Hijacker; it used this technique to hide its code and continually re-infect users. The infection also installed a Browser Helper Object (BHO) — a small program that plugs into your Web browser and directs its behavior. The name of the hidden CWS DLL file in the `AppInit_DLL`'s value was invisible to `Regedit`, and could only be "seen" by using an alternative Registry editor called `Registrar Lite`,

or by examining the exported Windows Registry key. The fact that the DLL was difficult to see made it a lot harder to remove. How can you remove a file you don't know the name of? The CWS-hidden DLL also changed its name after several reboots or after botched removal attempts — so even if it was successfully identified, a reboot it might cause it to adopt a completely new name. This practice of using morphing file names has since been adopted by several rootkits, including `Apropos`, `Gromozon`, and other stubborn malware including the `Vundo` trojan (which has a rootkit variant).

Windows did not intend for the `AppInit_DLLs` feature to be used for malicious purposes. Some very useful programs such as antivirus software and software debuggers use `AppInit_DLLs` injection constructively. In fact, the `AppInit_DLLs` value was put to very good use by the security researcher, Ilfak Guilfanov, when he developed the WMF Metafile HotFix, which utilized `AppInit_DLLs` injection to patch the zero day WMF Metafile vulnerability before the Microsoft patch was released.

Privileged hooks

In Windows NT systems (Windows NT 3.1 and later), an application program can operate in kernel mode by using a kernel-level *device driver* — a small program (implemented by launching a Windows service) that gives user programs unrestricted access to memory, hardware, and CPU privileged instructions. To a rootkit, that kind of access is as irresistible as a pirate's treasure chest — but the only way to get it is to use *privileged hooks* — which operate much like unprivileged hooks (described earlier in this chapter), with one major difference: Privileged hooks alter the data structures or tables used by the kernel itself. Because all programs and processes need to access kernel functions, modifying a kernel data structure affects all running programs.



NT stands for *New Technology*, and the term *NT-based Windows systems* encompasses Microsoft's family of 32-bit multitasking Windows operating systems, the first of which was Windows NT 3.1 introduced in late 1992, Windows NT 3.x and 4.0 Workstation and Server versions, Windows 2000, Windows XP, Windows Server 2003, and now Windows Vista.

Installing drivers as rootkits

Installing a driver is a legal route that user programs take to get into the kernel, but it is also one that blackhat rootkit authors have maliciously exploited to infect computers with rootkits.



A driver doesn't have to actually represent a physical device (though many do); to a rootkit, it's a method of obtaining kernel-level privileges rather than a means to operate a specific piece of hardware such as a printer or mouse.

By definition, device drivers operate in kernel mode. They have privileges at Ring 0 — as high on the food chain of privileges as you can get — and that means only a privileged hook can bend the kernel's functions to suit a rootkit's purpose. So a kernel-mode rootkit gains control of the kernel by using a device driver that runs as a Windows service. You're probably familiar with many of the *services* that come preinstalled with the Windows

operating system. These processes run in the background and accomplish many tasks behind the scenes — such as Windows Automatic Updating and Event Logging.



To see a list of Windows services, you can access the Services Console by choosing Start ⇨ Run, typing **services.msc** into the Open: field, and pressing Enter. You will notice that there are both Microsoft and application program services listed. Naturally, any fully functional rootkits that might exist in your system won't be listed; they've made sure that the Services Console will ignore them.

Why malware writers use device drivers

Malware writers have determined (accurately) that loading a kernel-mode driver into a system gives them nearly unlimited access. No wonder the preferred rootkit methodology is to load a device driver so the rootkit can intercept and redirect native API calls — and thereby exert a global impact on the entire system. A kernel-mode rootkit can overwrite the basic data structures that the operating system uses to assess the system status — so that those structures fail to report anything that may expose not only the rootkit but also its malware-related processes, files, directories, ports, and Registry entries. Result: The whole infected operation is hidden to a degree that user-mode rootkits rarely match. (Not to mention that user-mode rootkits can easily be ferreted out by rootkit detectors running in kernel mode.)

How installing a device driver allows easy kernel access

The Windows OS treats device drivers as services — and requires verification of digital signatures for device drivers to confirm that a driver is not only compatible with Windows but also hasn't changed since it was last tested. If either one of these conditions isn't met, Windows triggers a user alert — but Windows leaves it up to the users to decide how they want to handle an unsigned driver. For enhanced security, 64-bit versions of Windows Vista depart from this policy and prevents any unsigned device drivers from loading. This is one of the security features implemented in 64-bit Vista to keep rootkits out of the kernel.

Currently, rootkit coders can bypass Windows driver verification by writing a program that relies on low-level API calls to install their driver. This suppresses the alerts that are normally triggered when a program loads an unsigned driver. Pretty sneaky!

Developing a kernel device driver is accomplished by creating a user-mode application to complete that task. Microsoft actually supplies a software-development kit (SDK) to help legitimate programmers who are interested in coding kernel device drivers. Three consistencies make this process easier:

- ✓ Kernel device drivers always follow the same naming convention; they all have a `.SYS` file extension. For example, a device driver might have a name like `mydriver.sys`.
- ✓ A file with a `.DLL` extension differs from a `.SYS` file in only one respect: the specific DLL files it links to.
- ✓ `SYS` files always link to `ntoskrnl.exe` (short for *NT operating system kernel*) and often link to `hal.dll`; Win32 Portable Executable (PE) DLL and EXE files link to `kernel32.dll` and `ntdll.dll`. Win32 PEs may also link to other system drivers depending on the type of support that's required for the Windows platform being used (some examples: `Ndis.sys` provides network driver support and `WmiLib.sys` provides support for Windows Management Instrumentation).

Registering a driver

After a driver is coded, it must be registered as a system service and then started to complete the installation process. Of course, coding the driver source code is the really complicated part. If that's done improperly, it can crash the system — and the infamous *blue screen of death* (BSOD) is a telltale sign of the system that many kernel-mode-rootkit-infected computers experience. It stems from the difficulty involved in coding a kernel device driver.

More about drivers

If you're curious to learn more about the process of developing a kernel device driver, refer to these links:

- ✓ Windows NT Kernel Programming — Introduction to Device Drivers:

www.catch22.net/tuts/kernel1.asp

- ✓ Instructions on creating DLLs:

www.thevbzone.com/l_dll.htm

Hooking the System Service Descriptor Table (SSDT hooking)

Like user programs, the kernel also maintains its own set of tables that it uses to carry out system functions, interact with hardware devices, and manage its own resources. In fact, every system call from a user program (via a user-level API) is eventually translated into an entry in a data structure that the kernel uses to locate the functions it needs to perform — the System Service Descriptor Table (SSDT). The kernel is where Windows actually performs the work that completes user-program requests; the kernel uses the SSDT like the index in a book to efficiently find the instructions it must perform. Each entry in the SSDT points to a set of instructions that performs a requested user program function. The functions themselves are called system services, and they are located in the NT operating system kernel, aka `ntoskrnl.exe`.

Now, kernel-mode rootkits alter kernel data structures or tables — and most of them hook the SSDT. Because all system calls eventually funnel down to functions that the SSDT points to, modifying the entries in this one kernel data structure enables a rootkit to affect *all* user programs that access it. It's a very powerful and efficient approach; — and if it's done correctly, the user should have no idea that the rootkit is there.

The SSDT stores the entry-level addresses of kernel APIs — and being able to modify the pointers in this table is like winning the lottery to a rootkit. No fiddling around with pointers to user-mode APIs here; kernel-level control is the goal.

A rootkit hooks the SSDT by overwriting the addresses in that table with pointers to rootkit functions, instead of legitimate system services — a process known as *kernel patching* or *native API hooking*. Then, when a call is made to execute a system service that the rootkit has patched, the service dispatcher retrieves and invokes the service that the patched pointer directs it to — a which just happens to be a rootkit function (located within the rootkit driver), of course. Because patching the SSDT affects all user processes that reference that structure, it is commonly referred to as installing a *system-wide hook* or a *global hook*.



Installing a global hook is not only the domain of rootkit programs but many security programs also go about their business by installing global hooks to intercept and monitor key kernel APIs — so it is important to identify the source of the global hook when attempting to decipher security program alerts that notify you of a global hook installation.

Though many kernel-mode rootkits function by hooking the pointers to native APIs in the SSDT, this is an older technique that is pretty much universally detected by anti-rootkit programs. So a more sophisticated technique, known as *direct kernel-object modification* (DKOM) has become the new, preferred methodology for advanced rootkits (more about that in the “Direct kernel-object manipulation” section later in the chapter).



There are over 200 Native APIs in `ntdll.dll` and most of them are undocumented. They all begin with an `Nt` prefix, but when these same APIs are exported, the names of the functions change — an `Nt` prefix is swapped for a `Zw` prefix— so each single `ntdll.dll` function is referred to by two names. For example, the native API in `ntdll.dll` that opens or creates a Registry key is called `NtCreateKey`, but when that same function is exported by `ntdll.dll`, it is called `ZwCreateKey`. The Metasploit (a Web site that provides resources to intrusion-testing researchers) maintains a list of the native API entries indexed by the SSDT (aka the System Call Table). You can find it at the following URL:

```
www.metasploit.com/users/opcode/syscalls.html
```

Kernel inline hooking

Both user-mode and kernel-mode rootkits can actually overwrite user-mode API and kernel-function code, respectively, via a technique called *inline function hooking*. We've already described how user-mode inline hooking works earlier in the chapter. Kernel-mode rootkits also use inline function hooking — but kernel-mode inline hooking involves altering the code that implements system services in the kernel (`Ntoskrnl.exe`). The rootkit hooks a system service by adding a jump instruction to the system service code (pointed to by the SSDT). When the jump instruction is executed, it immediately redirects processing to a rootkit-supplied function. What happens next varies, depending on what the rootkit code says to do — when the rootkit code is executed, it may duplicate the original legitimate kernel code but filter the results — or it may pass control back to the original legitimate kernel function, allow it to complete its work satisfactorily, and then bounce control back to the rootkit function, which then filters the results. Either way, the primary objective is to allow the function to process normally while excluding evidence of the rootkit. Blackhats tend to prefer this technique over hooking the SSDT or IDT (discussed next) because inline function hooking at the kernel level is harder to detect than SSDT hooking — but it's also harder to implement.



There are a few rootkit detectors that can detect kernel inline function hooking, and two of the programs that are able to do that, namely Gmer and DarkSpy, are included on the CD. (They are discussed in Chapter 9 and in the Appendix.)

SYSENTER hooking

SYSENTER hooking is so named because it hooks the entry point to the SSDT — the crossover from user mode to kernel mode. By hooking this juncture, a rootkit can filter all system calls *without* patching the SSDT. A rootkit does this by overwriting the pointer to the System Service Dispatcher with a pointer to code in its own driver. This way, whenever a user-level API (and subsequent SYSENTER) is called, the rootkit decides whether it should direct execution to a replacement function (to stay hidden) or let the real system service in the kernel handle it. Probably, because so few anti-rootkits programs were able to detect it, SYSENTER hooking was recently used very successfully by a new rootkit dubbed pe386 (more about pe386 in Chapter 12).

Replacing the SSDT

Another interception technique is a bit more drastic: to actually *replace* the SSDT and redirect all system calls to the “new” SSDT, which is actually a copy of the real SSDT with additional pointers to new system services appended. When this approach is used by a rootkit, as you might expect, the new additions would reference rootkit services (in the rootkit driver) rather than legitimate kernel services.



A variation of this technique is used in a beneficial way by some antivirus providers to monitor and intercept actions that could pose a threat to user systems.

Hooking the Interrupt Descriptor Table (IDT)

Kernel-mode rootkits can also modify the Interrupt Descriptor Table to alter the execution of programs. This technique works because an *interrupt* is generated when either a hardware device or program requests CPU time while another program is executing; they have to take turns. For example, if you've scheduled an anti-spyware scan for a specific time, a software interrupt may be issued when that time arrives. A printer may issue a hardware interrupt so it can execute a print job.

Naturally the OS has to keep track of what it's doing and what it's interrupting; the Interrupt Descriptor Table is a system data structure that contains pointers to *Interrupt service routines* (ISRs). Rootkits can replace the pointers to ISR with pointers to their own code.

When an interrupt is triggered, program flow is returned to the operating system, which then executes an appropriate Interrupt service routine — and we're back to services, a known rootkit target (as described earlier in this chapter). If the Interrupt Descriptor Table has been hooked and doctored, then an interrupt triggers and executes a rootkit routine instead of a normal ISR.

Using Even More Insidious Techniques to Hide Rootkits

Technology improves by its nature; rootkits are no exception. Besides hooking, kernel-mode rootkits have newer and stealthier alternative technology with which they can avoid detection — in particular, direct kernel-object manipulation (DKOM). The following sections describe this alternative approach to rootkit technology.

Direct kernel-object manipulation

Since kernel-mode hooking can be detected by several rootkit-detection programs, a newer — and stealthier — alternative to hooking is being used in the most advanced kernel-mode rootkits. This new technique is called *direct kernel-object manipulation* or DKOM. As we've discussed already, obtaining administrative privileges is paramount to a rootkit. Running processes inherit

their privileges from the user accounts they are associated with. DKOM can circumvent limited user privileges by selectively elevating the tokens (security objects that determine privilege) associated with rootkit processes (or with any process, for that matter) so they are equivalent to those of the administrator. That comes in handy if your aim is disabling a firewall, disabling anti-malware scanners, or installing a driver. DKOM even lets a rootkit process evade standard system monitors such as the Event Viewer by changing the group that process is identified with — a rootkit process can inconspicuously blend in by appearing belong to the system. No wonder the most sophisticated kernel rootkits forego hooking altogether and opt to use DKOM — or combine the two — to achieve their stealth.

DKOM enables kernel-mode rootkits to modify kernel structures that exist in memory — such as lists of active processes, of *threads* (individual process components), and of loaded drivers. Normally the kernel uses such lists to keep track of the information it must process — and they present additional opportunities for rootkit deception. Advanced rootkits can use DKOM to manipulate these lists, erasing the traces of untoward activity.

Here are some signature characteristics of DKOM:

- ✓ DKOM is different from hooking because the kernel data structures themselves are modified, not just table entries. Many advanced kernel-mode rootkits use this technique to hide active processes and drivers more effectively from rootkit-detection programs.
- ✓ DKOM *can* prevent you from seeing a suspicious rootkit process because you won't even know the rootkit process is executing. Processes can also be hidden by hooking APIs (as described earlier in the chapter), but even kernel hooking is a technique that most anti-rootkit programs can detect very easily (with the exception of inline function hooking). DKOM can very successfully hide rootkit processes from many rootkit detectors that are normally capable of detecting them.
- ✓ One of the objects manipulated by DKOM, is a data structure the kernel maintains called the *Process List*. The Process List is composed of a linked list of Executive Process Blocks or EPROCESS blocks. Each EPROCESS block represents an active process in memory. An EPROCESS block contains unique identifying information about its respective process, including a pointer to the previous process and a pointer to the next process, in the list. The pointers allow the list to be read out of sequence — by following the pointer to the “next process” in each EPROCESS block. A rootkit using DKOM can manipulate these pointers so a rootkit's process ERPROCESS block is unlinked from the list. This effectively makes Windows skip the rootkit process when assembling the list of active processes. Suddenly the rootkit process “isn't there.”

Figure 7-1 illustrates how a rootkit uses DKOM to hide a process from Task Manager or even an anti-rootkit program incapable of detecting DKOM of the Process List. By rearranging the pointers in the Process List, the rootkit EPROCESS link is excluded from the list of active processes.




Figure 7-1:
Using
DKOM to
hide a
process.

Now here's a strange twist that works in a rootkit's favor — as long as a process is already active, removing a process from the Process List doesn't prevent it from executing. (Devious, isn't it?) That's because another data structure — the process *thread* list — is the one that is actually consulted for process scheduling. Processes are executed by running the individual components that the processes are composed of — its threads — and a single process can contain multiple threads. However, the Process List is what Task Manager and other Process viewers refer to when they assemble the list of active processes. If a rootkit manipulates the Process List to exclude its own processes, then its processes may appear to be invisible — but they will still be executing behind the scenes.



Klister, a utility developed by Joanna Rutkowska, can detect rootkits that use DKOM to hide their processes from the Process List. Klister can be downloaded here:

www.invisiblethings.org/tools/klister-0.4.zip

Trojanized utilities

The earliest UNIX rootkits were much simpler than today's super-stealthy versions, and what their authors did was replace system files (utilities) with counterfeit copies of their own. These new (trojanized) utilities were designed to perform their normal functions while secretly hiding the rootkit. To use a parallel example relevant to Windows users, rootkits might replace Task Manager's executable file (`taskmgr.exe`) so that all active processes except those belonging to the rootkit were listed.

Direct replacements of system utilities are easily detectable (as we shall see in Chapter 9) and pretty easy to reverse by installing a new copy of the affected executable. Today, malware writers may try to replace an authentic file with their own, but this is difficult to do to system files because of Windows File Protection, which we discuss earlier in this chapter. Even if Windows File Protection could be bypassed, then an antivirus or anti-trojan should be able to pick up the bogus files in a routine scan. Another obstacle to replacing critical system files with a trojanized version is the difficulty of maintaining the original functionality of the replaced file. Since Windows is not an open-source operating system, all of its code is proprietary. It is not an easy task to write *rewrite* the code of an essential system file so it incorporates both normal *and* rootkit functionality — doing so might cause a system crash.

Even so, user-mode rootkits may make changes to critical system files by replacing programs with trojanized utilities of their own, or having the files depend on trojan DLLs that run within the context of a system process. A process called *checksum analysis* can help you discover trojanized utilities; see Chapter 9 for more information on checksum analysis.

Looking into the Shady Future of Rootkits

Some rootkits are actually on the side of the white hats — for now, anyway — because they are original laboratory creations, often designed by rootkit researchers for beneficial purposes. They’re categorized as non-public or (well, yeah) laboratory-based — which means they’ve been developed as research tools and used to gain insight into new techniques that future rootkits could possibly use to hide. Such proof-of-concept rootkits are used to develop tools or strategies to combat potential techniques before a malicious rootkit starts using them in the wild. Three non-public rootkits that fall under this classification are Shadow Walker, FUTo, and Subvirt. Here’s a look at how these three rootkits incorporate new hiding techniques — and how they’ve taken rootkit detectors to a whole new level.

Hiding processes by doctoring the PspCidTable

An even newer method of process hiding has been built into the laboratory rootkit called FUTo. FUTo was christened so by its author, Peter Silberman, because it is the successor to the publicly available rootkit FU (or FU II, which became FUTo). Even though most of these laboratory rootkits were developed to challenge and improve anti-rootkit tools, many are available for download on the Web, which means tweaked versions may find their way into distributed malware. Consequently, a version of FU is used in the “Aim Virus” (a misnomer) rootkit, the Fanbot worm, and several spyware/adware programs — proving you don’t have to be a kernel programmer to bind FU to a threat.

Besides hiding processes by DKOM of the Process List, FUTo alters the only data structure used by the operating system to keep track of both active processes *and* threads — PspCidTable. If a rootkit hides a process using DKOM of the Process List both IceSword and BlackLight are able to detect it, but neither can successfully detect FUTo’s added method of process hiding.

Until FUTO made an appearance, restoring the Process List to its original state was sufficient to reveal hidden rootkit processes. No more; FUTO introduced a new approach to hiding processes by modifying the PspCidTable; it also created a need for a detection utility that can target that technique. That tool is RAIDE, and we discuss it in Chapter 9.



DarkSpy and the GMER rootkit detection programs included on the CD are capable of detecting both FU's and FUTO's method of hiding processes and drivers.



A very simple and quick freeware utility called kproccheck can detect rootkits that hide their processes using both of FUTO's methods. Kproccheck is one of three very effective command line tools developed by the Security Information Integrity Group (SI^{IG}), and we will be talking about another one of their programs (AntiHookExec) in Chapter 9. Kproccheck may be downloaded here:

www.security.org.sg/code/kproccheck.html

Hooking the virtual memory manager

A new laboratory creature uses virtualization to make itself invisible to any detection strategies employed from within the target computer. *Virtualization* is a general technique used to create a separate simulated environment (complete with operating system) within the host computer. A virtual machine is like a "machine within a machine," and it behaves and functions separately and independently. A program launched within a virtual machine is normally only aware of the virtual environment that surrounds it, but it views that world as the entire computer.

This new rootkit technique uses virtualization to install itself at a level below that of the host operating system. Because it sits at a level that is lower than the host computer's real operating system (effectively right above the hardware), it is able to see not only its own environment, but everything that happens on the computer. However, booting from an external operating system can successfully detect this virtual-machine-based rootkit and its associated malware because an external operating system is able to see the untainted view of the host machine, including the virtual-machine-based rootkit. Shadow Walker is a non-public, in-memory rootkit that works by hooking the Virtual Memory Manager (VMM). ShadowWalker is discussed in detail in Chapter 12.

Virtual-machine-based rootkits

Another advanced rootkit technique developed in the laboratory is the virtual-machine-based rootkit (VMBR). Microsoft and the University of Michigan recently coauthored a VMBR that is essentially undetectable by any method utilized from within the target computer. The project is labeled SubVirt, which is a moniker derived by combining portions of the words *subvert* and *virtual*. The term *subvert* is often used to describe a rootkit's ability to trick the operating system into believing it doesn't exist, which is why Greg Hogland and Jamie Butler decided to call their book *Rootkits: Subverting the Windows Kernel* — and also why we have tried to avoid using that word too often in this book.

SubVirt installs and runs a virtual machine monitor (VMM) in its own area of disk space that is totally undetectable and off-limits to the host operating system. This VMM behaves like a self-sufficient machine (complete with its own operating system) that controls and keeps tabs on what's going on the entire computer — hence the term *monitor*.

Yet, because it effectively operates at a level below the host kernel, SubVirt still has access to all the layers above it — but remains totally inaccessible to the host operating system. The rootkit also places the original host operating system inside a virtual machine. After the host operating system is isolated in this way, its security tools become totally ineffectual against the VMBR. That means any malware that the VMBR installs inside the virtual machine monitor — or any newly spawned Virtual Machines the VMBR creates to further compromise the infected computer — are also invisible to the host.

But even this creature has a weakness: The VMBR must find and exploit a host vulnerability that will enable it to alter the master boot sequence. The VMBR has to load the virtual machine *before* the host operating system loads, so it can sit at a level that underlies the host's operating system. The laboratory VMBR experiment assumes the presence of such host vulnerability (a premise that critics of SubVirt claim is quite difficult to implement).

In the lab, SubVirt successfully installed itself and ran four malicious services that targeted host resources. Since the SubVirt exists at a level in between the hardware and the host operating system, an appropriate countermeasure has been proposed: incorporating detection into a hardware chip that enables the host computer access to the virtual machine files.

An easier disinfectant approach would be to boot from an external operating system contained on an alternate medium (such as a CD-ROM or USB flash drive). The external operating system would offer an uncompromised view

of the *entire* target computer. As such, it would present an effective way to detect and remove malware — using conventional scanners resident on the external medium.

The SubVirt rootkit was successfully installed and run on a Linux platform using VMWare, and on a Windows platform using Microsoft's VirtualPC emulation software.



It is important to emphasize that a VMBR is not about exploiting virtual machine vulnerabilities. Using VMWare or VirtualPC does not make you a more likely target for a VMBR; it actually makes you less vulnerable to VMBR exploits (as discussed in Chapter 6). An operating system that's already running within a user-installed virtual machine would always operate at a level below a VMBR — which would render it less prone to exploit than the normal physical operating system. That's because the lowest layer can see everything above it — while remaining invisible to all the layers above it. Whatever entity establishes itself at the lowest layer — malware or the host operating system — has a distinct advantage and occupies a position of control.

Some critics have described virtual-machine-based rootkits as impractical to implement in the real world. One reason for that is that virtual machines place quite a drain on system resources — particularly memory, which would be difficult for a user to miss. The disk space required is also a consideration, since SubVirt comes with its own operating system in tow. Still, as hardware and memory advance, this may not always be the case. Complexity of installation is another factor that may be prohibitive. Even so, the SubVirt authors contend that their model presents a viable future threat. Without a doubt, virtual-machine-based rootkits would be extremely dangerous, were they to become a reality outside the lab.

Chapter 8

Sniffing Out Rootkits

In This Chapter

- ▶ Watching your network for signs of rootkits
 - ▶ Getting to the bottom of lockups and other odd behavior
 - ▶ Checking your system resources
-

One of the last things anyone wants to hear, be they network administrators or individuals with standalone computers, is that a backdoor or rootkit has been sneakily installed on their systems. Even the possibility of such a condition merits thorough investigation — and perhaps painful acknowledgment.

This chapter, along with Chapter 9, shows you how to discover the presence of a backdoor trojan or a rootkit on your standalone system or network. A similar investigative process applies to either process — albeit on a smaller scale for an individual computer than for a network. This chapter shows you how to search for and find rootkits on a network — but much of what you find here applies to standalone computers as well.

Watching Your Network for Signs of Rootkits

Although your network may be well fortified, it's still wise to monitor network activity for signs of a rootkit. Most rootkits are installed on a computer to hide illicit backdoor activity. Just as a seemingly legitimate business can be a front for an illegal one — or the Web can provide the anonymity that allows a cyber-criminal to thrive — your network's bustling activity can provide a comfortable milieu for a rootkit.

Fortunately, some key indicators of rootkit activity can be monitored. The first step in sniffing out a rootkit on your network is to establish whether there's a likelihood of infection by an intruder. You do that by detecting whether any signs of intrusion exist.



We've said it before but it bears repeating: You should never become complacent about network security. The current threat landscape demands that you always be on your toes. Intruders are always seeking out vulnerabilities; if you let your guard down, even for a moment, your computer or network can become a blackhat hacker's target.

Watching logs for clues

It's true that rootkits can intercept operating-system APIs and get them to alter nearly anything the system reports. Yet they may not totally succeed in hiding themselves; at least the log entries that were tied to the initial break-in attempt should remain intact — and visible upon inspection. Here's a list of logs to inspect (whether manually or automatically) to gain relevant clues:

- ✓ **Event logs:** Examine these for unsuccessful logon attempts. If the same party (as determined by IP address) that was trying to log on unsuccessfully suddenly logs on successfully — and an escalation in privileges from User to Administrator follows — then you've got trouble. Event logs should be examined for new processes and service installations. Typical suspicious doings include the creation of new user accounts, sudden upgrading of user privileges, and turning off (disabling) security programs — better investigate 'em.

These days, it's common for attackers to disable antivirus and firewall programs as part of their method of intrusion; if that happens, consider it a red flag. (To get a handle on whether this is happening to your system, see the “Investigating Lockups and Other Odd Behavior” section later in this chapter.)

- ✓ **Firewall logs:** Examine these for both inbound and outbound access attempts. Programs that “phone home” should be thoroughly investigated — up to and including the deciphering and tracing of IP addresses. Unusual trends and other anomalies should be noted and investigated. For example, take note if any processes that normally don't have Internet access start making outbound access attempts — especially if those attempts are repeated.





Some legitimate processes that are often targeted this kind of compromise are `Explorer.exe` and `WinLogon.exe`. That's because these system processes are always running from the time Windows starts — and they are the parent processes of other user and system processes. Connection attempts associated with such processes can indicate that a malicious DLL or thread is trying to hitch a ride by running within them.

If you normally use an alternative browser such as Firefox or Opera, look to see whether Internet Explorer is showing background Internet activity. Many suspect programs default to using Internet Explorer — and such activity can be a significant clue. Here are some guidelines:

- Be sure to confirm that the activity isn't due to the Windows Update process (`wuauclt.exe`).
- In general, any repetitive, blocked accesses should be investigated.
- Investigate when any of the programs that have approved Internet access show a sudden rash of unexplained Internet activity; make sure they're on the up and up. Use a *port-to-process mapping program* such as Port Explorer or TCPView to map open ports to the processes that own them.
- Determine the underlying processes of any outbound attempts by `svchost` to access the Internet. A program such as Port Explorer or TCPView can do that job — in combination with Process Explorer to trace process IDs (PIDs).
- All port-to-process mapping programs produce logs you can save for future comparison. If you're concerned about suspicious network activity, you can upload the logs to one of the security sites listed in Chapter 13 to get help with analyzing them. Gibson Research also offers free online leak tests for firewalls at www.grc.com/lt/leaktest.htm.
- See the sections “Defending your ports,” “Catching rootkits phoning home,” and “Examining Active Processes” later in this chapter for more information.

✓ **Sniffer logs:** In much the same way that Port Explorer can capture RAT (remote-access trojan) packets, a sniffer can capture all the traffic flowing across a network. If network monitoring has alerted you to potentially suspicious activity, you can use a sniffer to capture the data going through an unusually active port — and then log it to disk and analyze the sniffer log for suspicious entries. Programs such as Ethereal (described later in this chapter) are designed for this purpose. If you want to zero in on the activity associated with a particular port or

process, then Port Explorer is an excellent choice, especially for smaller networks and standalone systems. (See the “Trusting Sniffers and Firewalls to See What Windows Can’t” section later in this chapter for more information.)

- ✓ **Intrusion-Detection System (IDS) logs:** IDS systems produce logs that can be analyzed for break-in attempts. While a firewall acts as a gatekeeper for network traffic, either allowing or denying entry, an intrusion-detection system goes beyond this — examining network activity for suspicious patterns and behavior.

An IDS is absolutely essential for a network — in which case it’s called (well, yeah) a *network IDS* or *NIDS* — essentially a network watchdog that monitors and analyzes all network activity, applying rules (established indicators of suspicious behavior) to collected network data, looking for any untoward activity that suggests a possible break-in. For example, blackhat hackers use automatic tools to poke around on the Web looking for computers with open ports that are vulnerable to attack — a technique called *port probing* — which shows up in an NIDS log as a rash of TCP inbound access attempts across a wide range of ports on the network. If the NIDS monitors outbound traffic for suspicious behavior and sends an alert if it finds something out of the ordinary — but won’t stop it — then it’s considered *passive* rather than reactive protection. (For an example of a NIDS that offers a best-of-both-worlds approach to a typical port attack, see the nearby “Protecting ports with Snort” sidebar.)



A useful online resource dedicated to preventing Web-based attacks is McAfee’s HackerWatch. It offers free on-site testing for system vulnerabilities (such as open ports) and provides immediate feedback on whether vulnerabilities are detected on your system.

- HackerWatch can be found at www.hackerwatch.org/probe.
- Gibson Research offers free online testing for port vulnerability at www.grc.com/x/ne.dll?bh0bkyd2.

Protecting ports with Snort

One of the most popular and widely used NIDS products is Snort — a freeware utility that combines detection and prevention (which is probably why it’s so popular). Snort relies on a set of rules — frequently updated — to formulate and improve its algorithm (formula) for spotting

suspicious behavior. Snort also uses definitions for malware detection. You can download Snort from www.snort.org. (For the full story on Snort, pick up *Snort For Dummies* by Charlie Scott, Paul Wolfe, and Bert Hayes [Wiley Publishing, Inc., 2004].)



✓ **Host-Intrusion Prevention Systems (HIPS)** HIPS are proactive programs similar to their IDS brethren; they react to behavior that's symptomatic of intrusions — and block it. There are many HIPS programs available for both standalone and network computers. Some of the more popular ones are AntiHook an excellent freeware utility: (www.infoprocess.com.au/antihook.php), ProcessGuard by Diamond CS (www.diamondcs.com.au/processguard), and System Safety Monitor (<http://sysssafety.com>). Another popular alternative called Prevxl (www.prevx.com) combines malware scanning and HIPS into a single package. Prevxl can be used right out of the box, making it a perfect solution for beginners.

One of the best resources for understanding and comparing the features of different HIPS programs is an article called *HIPS/IDP Programs/Services* that can be found at CastleCops, a Web site dedicated to educating computer users of all levels:

```
wiki.castlecops.com/HIPS/IDP_programs/services
```

The article compares HIPS program offerings in a table and explains each HIPS feature. Clicking a program name in the table takes you to a summary of that program, a link to its vendor's Web site, and a support forum (if one exists). A more advanced article evaluates HIPS programs by conducting comprehensive testing; you can find it here:

```
http://kareldjag.over-blog.com/categorie-69553.html
```

Defending your ports

As mentioned elsewhere in the book, the most insidious rootkits install a *backdoor* (an unauthorized entry point) on the target so they have access to whatever your computer or network has to offer. A backdoor allows a remote attacker to gain control and direct server activity. It relies on information transfer between hacker and host — and this transfer can often be the means by which a rootkit is recognized. One big hint: A remote-access trojan (RAT) on an infected server communicates with its remote client through an open port or socket.

The following sections describe the ports that are available to your computer, explains how they can be hijacked by hackers, and shows you how you can monitor them — to identify potential illicit activity.

Knowing your ports

There are 65,535 TCP ports and 65,535 UDP ports available to a computer; the application that uses a port defines whether it's UDP or TCP. UDP allows rapid transfer of data with the trade-off of possible occasional packet (small

individual units that transmitted data is comprised of) loss. TCP transfer is less concerned with quantity than quality — and more concerned with reliability and accuracy so all data arrives at its final destination with minimal packet loss.

The following is a sampling of a few reserved ports and the specific server functions they supply:

- ✓ Ports in the 0 through 1024 range are reserved for certain server functions.
- ✓ Port 20 is used by File Transfer Protocol (FTP) to transfer files.
- ✓ Port 21 is used for FTP server and client file transfer or communication.
- ✓ Port 25 is used by the Simple Mail Transfer Protocol or SMTP e-mail services.
- ✓ Port 53 is used by your Domain Name Server (DNS) to map numerical IP addresses to the text names that are easily read and remembered like Google.com.



TIP

If you only have a numerical IP and want to translate it to a human-readable domain name, you can plug it into a reverse DNS database such as the one at <http://remote.12dt.com>. You can also use such *reverse DNS lookups* to determine the Internet service provider (ISP) for any numerical Internet address. To obtain the contact/registration information for a particular IP address or domain name, you can plug it into the Whois database at <http://whois.com> — and what you'll get back is the name of the network that owns the block of IP addresses that it belongs to. If you need a one-stop shop for everything IP- and DNS-related, try www.dnsstuff.com. It offers Whois and reverse DNS lookups under one roof, as well as a variety of other services.

- ✓ Port 80 is used for Hypertext Transfer Protocol (HTTP) — the Web protocol that enables data transfer between Web servers and the computers that connect to them on the Internet.
- ✓ Port 110 is used for Post Office Protocol (POP3) or POP e-mail services.
- ✓ Ports 1024 through 49151 are available for use by applications through a registration process that's administered by IANA.



TECHNICAL STUFF

Port assignments are coordinated by the Internet Assigned Number Authority (or IANA) — the same agency that coordinates the assignment of global IP addresses to specific domain names. *Reserved ports* are not available for assignment to user programs or applications. Also, Port assignment doesn't confer credibility to the application that uses the

port. Malicious programs can use particular ports just as easily as assigned programs — and malware distributors are (shall we say) not known for paying much attention to the rules except to violate them.



If you're interested in learning more about port assignments, Network Ice maintains a Port Knowledgebase that lists a slew of frequently seen TCP and UDP ports — and what they mean. Here's where to find it:

```
www.iss.net/security\_center/advice/Exploits/Ports/
```

How blackhat hackers use ports

Hackers look for system vulnerabilities they can exploit. They can take advantage of an unpatched system or they can use social-engineering tactics that invite the downloading of bogus “helpful programs” or e-mail attachments. After a threat is installed on a system, it often attempts to establish an Internet connection and open a listening port. Smart trojans try to get a free ride by using a recognizable reserved port that they know will be acceptable to most firewalls. For example, the majority of network firewalls are configured to accept connections for Port 53 — the port used by your Domain Name Server (DNS). If a RAT uses Port 53, it can usually set up shop without triggering a firewall alert. Quite a few backdoor threats — many with stealth characteristics — use this port. Other reserved ports can be exploited in much the same way. For example, Port 25 — used normally for e-mail transmission via Simple Mail Transfer Protocol (SMTP) — is a favorite target for many trojans.

Trojans can use nearly any port — but some are associated with dedicated or specific ports — and (as you might expect) they don't comply with IANA guidelines — so it's possible to identify some RATs by the ports they use. Of course, just because you have evidence of specific port usage doesn't necessarily mean you have a backdoor trojan infection — or, if you do have one, that it's using *only* the port you've identified. Some trojans maintain *two* ports — one for listening to their controller commands and another for transferring information. (The Netbus trojan described in Chapter 9 maintains two open ports for this purpose.)

Two very prevalent rootkits — Hacker Defender and Win32/Hackdoor.B — hijack ports that are already in use by legitimate applications. As explained earlier, trojans can easily bypass your firewall by injecting DLLs or threads into processes that your firewall has already given clearance to. Bottom line: Those who create backdoors and rootkits are adept at finding ways to bypass a firewall. Some just turn it off altogether — and then disable Security Center notification so the user doesn't even realize the firewall's turned off.

Here are a couple of links to online databases that list ports commonly used by trojans:

- ✔ **Simovits**, a security consulting group maintains a list of ports known to be used by trojans. Here's where to find this valuable resource:

```
www.simovits.com/trojans/trojans.html
```

- ✔ **Gibson Research Corporation**: Known for its security research and security-related products, Gibson offers an interactive Internet database where you can plug in any port to get information on its uses — and even probe it on your computer to see whether it's visible (vulnerable) or stealthed (invisible). It's no accident that this company crops up earlier in this chapter — and it features Nancy's favorite place to research anything port-related:

```
www.grc.com/PortDataHelp.htm
```

- ✔ **SANS**: The SANS Institute is a valuable resource for everything related to Internet security. In addition to a list of trojans' favorite ports, SANS maintains an Internet Storm Center that issues alerts and warnings on current threats. You can find Internet security information and port specifics online at the following URL:

```
www.sans.org/resources/idfaq/oddports.php
```

Checking ports with port-to-process mapping

You're in luck: Several programs exist that enable you to view the open ports on your system and trace them back to the programs or processes that opened them. This tracing analysis is known as *port-to-process mapping*. Port-to-process mapping utilities — when used in combination with a process database and a trojan-port database — can help you troubleshoot trojan intrusions.

In the following sections, we discuss five tools that can do port-to-process mapping. Here's a quick overview:

- ✔ Netstat is a basic utility that comes with Windows. You must run Netstat from the command line, and you can use different parameters to manipulate its output.
- ✔ TCPView by Sysinternals offers a graphical user interface (GUI) and additional features that make it a little easier to use than Netstat. TCPView produces an activity log that can be saved to a file, which is handy if you want to seek help in analyzing its output.



- ✔ IceSword is an anti-rootkit program that has a handy Port function. (More about IceSword in a minute.)
- ✔ Port Explorer by Diamond CS provides a high-end, colorful GUI and many extended features such as packet spying and testing ports against an onboard trojan-signature database.
- ✔ DarkSpy, an anti-rootkit program included on your DART CD, can display the ports hidden by rootkits.

Netstat

Netstat (short for *network statistics*) is a command-line tool that comes with most operating systems, including Windows. Netstat enables you to see open ports and (normally, anyway) the processes that have opened them. Its output is rather primitive but adequate — as long as you aren't up against a sophisticated rootkit. You invoke Netstat at the command line, in just two steps:

1. **Click Start, choose Run, type cmd into the Open: field, and press Enter.**
2. **Type the following command at the C:\ command prompt (after getting to the C: directory by typing cd\ and pressing Enter):**

```
netstat -a -b | more
```

This command performs three tasks:

- a shows all connections and ports.
- b shows the process that created this port.

| more *pipes* (redirects) the output so appears on-screen one screenful at a time; you “press any key to continue” to see the next screenful.

The command generates a list of processes and their associated ports, as illustrated in Figure 8-1. If you want to save the output to a text file (here we call it `myports.txt`), issue the following command instead:

```
netstat -a -b > myports.txt
```

Unfortunately, most rootkit ports are not visible to Netstat. That's because it's frequently targeted by rootkits because it's available to all Windows users and is simple to use. Still, Netstat is worth using anyway, at least in the early stages of your research when you have not yet determined whether your problems are due to ordinary malware or a rootkit. If a rootkit does show up later in the investigation, the fact that it hid from Netstat is an important piece of evidence that can be used to help determine whether a threat is a rootkit.

```

C:\WINDOWS\system32\cmd.exe
C:\>
C:\>netstat -a -b | more
Active Connections
Proto Local Address           Foreign Address         State       PID
TCP    SOPHIA:microsoft-ds     SOPHIA:0                LISTENING   4
[System]
TCP    SOPHIA:netbios-ssn     SOPHIA:0                LISTENING   4
[System]
TCP    SOPHIA:1281            localhost:1282          ESTABLISHED 988
[MOZILLA.ESE]
TCP    SOPHIA:1282            localhost:1281          ESTABLISHED 988
[MOZILLA.ESE]
UDP    SOPHIA:4500            ***                      804
[lsass.exe]
UDP    SOPHIA:1478            ***                      3380
[procexp.exe]
UDP    SOPHIA:isakmp           ***                      804
[lsass.exe]
UDP    SOPHIA:microsoft-ds     ***                      4
[System]
UDP    SOPHIA:3319            ***                      1284
C:\WINDOWS\system32\dcus2.dll
C:\WINDOWS\system32\lmon.dll
C:\WINDOWS\system32\ULDP32.dll
C:\WINDOWS\system32\winrmr.dll
c:\windows\system32\W32_32.dll
c:\windows\system32\w32time.dll
[svchost.exe]
UDP    SOPHIA:1401            ***                      1284
C:\WINDOWS\system32\dcus2.dll
C:\WINDOWS\system32\lmon.dll
C:\WINDOWS\system32\ULDP32.dll
C:\WINDOWS\system32\winrmr.dll
c:\windows\system32\W32_32.dll
c:\windows\system32\w32time.dll
[svchost.exe]
UDP    SOPHIA:3320            ***                      1284
C:\WINDOWS\system32\dcus2.dll
C:\WINDOWS\system32\lmon.dll
C:\WINDOWS\system32\ULDP32.dll
C:\WINDOWS\system32\winrmr.dll
c:\windows\system32\W32_32.dll
c:\windows\system32\w32time.dll
[svchost.exe]
UDP    SOPHIA:ntp              ***                      1284
C:\WINDOWS\system32\dcus2.dll

```

Figure 8-1:
The Netstat
list of
processes.

TCPView

TCPView is port to process mapper that has a GUI with a prettier display than Netstat. It's also easier to use and offers a few more display features. TCPView is freeware program offered by Sysinternals that can be downloaded from the following URL:

www.sysinternals.com/Utilities/TcpView.html

TCPView can display either a process view or an IP view (showing the remote IP to which a port is connected). It also enables you to save a log — important if you intend to seek troubleshooting assistance.

Figure 8-2 shows an illustration of TCPView's display.

One advantage of running TCPView is that you can start it from the run line (click Start and choose Run) in combination with AntiHookExec (as described in Chapter 9). If a rootkit has hooked some user-level APIs that allow port enumeration, this command restores the hooked APIs — making the rootkit ports appear on-screen. The syntax of the command looks like this if you launch it

from the run line (assuming that TCPView is in the C:\Program Files\TCPView\ folder):

```
AntiHookExec "C:\Program Files\TCPView\Tcpview.exe"
```




Figure 8-2:
What you
see with
TCPView
on-screen.



We cannot stress enough how valuable the Sysinternals programs are. Your DART CD includes two system analysis utilities: Autoruns (a program autostart viewer) and Process Explorer (a process viewer), not to mention a dedicated rootkit-detection program called Rootkit Revealer. You should make an effort to install and become familiar with these programs; each is an invaluable system-analysis tool in its own right. They will help you investigate areas of your system that malware is known to target. Systematic monitoring with these tools will enable you to spot anomalies before they get a stranglehold on your computer or network. You can launch them on the run line with AntiHookExec (as discussed in Chapter 9), enabling you to visualize user-mode rootkit components that would otherwise escape detection.

IceSword

IceSword is a rootkit-detection-and-removal program that works like a set of security tools for monitoring your computer. (We detail IceSword's many features in Chapter 9, and show how to get hold of it in Bonus Chapter 1.) You can access the Port function by clicking the Port icon on IceSword's Functions toolbar; the Port function lists all open ports — including hidden rootkit ports. Because IceSword does not distinguish *rooted* (rootkit-hidden)

ports from normal ones, it is best to compare its port listing to the output of another utility that cannot list rooted ports and see if there are any differences. That makes at least one security procedure relatively painless: Just run the program and inspect your ports on a routine basis. Figure 8-3 illustrates the Port display in IceSword.

Proto...	Local Address	Foreign Address	State	PID	PathName
TCP	127.0.0.1 : 1281	127.0.0.1 : 1282	ESTABLISHED	988	C:\PROGRA~1\mozilla.org\Mozilla\
TCP	192.168.1.101 : 139	0.0.0.0 : 0	LISTENING	4	NT OS Kernel
TCP	0.0.0.0 : 445	0.0.0.0 : 0	LISTENING	4	NT OS Kernel
UDP	127.0.0.1 : 3692	*.*		1284	C:\WINDOWS\SYSTEM32\svchost...
UDP	127.0.0.1 : 3320	*.*		1284	C:\WINDOWS\SYSTEM32\svchost...
UDP	127.0.0.1 : 1401	*.*		1284	C:\WINDOWS\SYSTEM32\svchost...
UDP	0.0.0.0 : 500	*.*		804	C:\WINDOWS\SYSTEM32\sess.exe
UDP	192.168.1.101 : 137	*.*		4	NT OS Kernel
UDP	0.0.0.0 : 1478	*.*		3380	C:\Program Files\Process Explorer\
UDP	192.168.1.101 : 138	*.*		4	NT OS Kernel
UDP	127.0.0.1 : 123	*.*		1284	C:\WINDOWS\SYSTEM32\svchost...
UDP	127.0.0.1 : 3500	*.*		1284	C:\WINDOWS\SYSTEM32\svchost...
UDP	127.0.0.1 : 4287	*.*		1284	C:\WINDOWS\SYSTEM32\svchost...
UDP	192.168.1.101 : 123	*.*		1284	C:\WINDOWS\SYSTEM32\svchost...
UDP	127.0.0.1 : 3691	*.*		1284	C:\WINDOWS\SYSTEM32\svchost...
UDP	0.0.0.0 : 4288	*.*		1284	C:\WINDOWS\SYSTEM32\svchost...
UDP	127.0.0.1 : 3501	*.*		1284	C:\WINDOWS\SYSTEM32\svchost...
UDP	127.0.0.1 : 1400	*.*		1284	C:\WINDOWS\SYSTEM32\svchost...
UDP	0.0.0.0 : 4500	*.*		804	C:\WINDOWS\SYSTEM32\sess.exe
UDP	127.0.0.1 : 3319	*.*		1284	C:\WINDOWS\SYSTEM32\svchost...
UDP	0.0.0.0 : 445	*.*		4	NT OS Kernel
RAW	---	---	---	4	NT OS Kernel
RAW	---	---	---	4	NT OS Kernel
RAW	---	---	---	4	NT OS Kernel

Figure 8-3:
The
IceSword
port
functions.

Dark Spy

DarkSpy is a very effective anti-rootkit program that offers numerous functions for viewing hidden rootkit components, much like IceSword. It's the only rootkit-detection tool on the CD with a port-viewing function that can expose rootkit-hidden ports, even though it can't tell you which process opened the port (IceSword does that, by the way). Though rootkit ports won't be distinguishable from others listed in the screen display (they're not color-coded), you can compare DarkSpy's port listing to that of Netstat (or any other conventional port-enumeration program we've mentioned) to zero-in on port discrepancies that may be rootkit-related.

We have included DarkSpy on the DART CD — running it (or the other anti-rootkit programs) from there should yield untainted and trustworthy results when you're looking for malware or monitoring ports. If you have a rootkit lurking, the programs on the CD should be able to detect it.

Port Explorer

Port Explorer is a port-to-process-mapping program developed by DiamondCS. It's available for a free download at this link:



www.diamondcs.com.au/portexplorer/

Port Explorer is the Rolls Royce of port-to-process mappers. Its plethora of features allow you to identify, monitor, and stop backdoor trojan processes — and effectively analyze any programs that attempt to phone home. (See the “Catching rootkits phoning home” section, later in this chapter, for details.) Using Port Explorer, you can kill trojan processes, disable data transfer across open ports, and still be able to analyze the content of the attempted transmission. This last feature is stellar; you don’t have to allow a malicious transmission to occur just to capture the data.

We consider Port Explorer worth every penny of its moderate one-time fee, should you decide to purchase it. Port Explorer has its own product-support forum — and an extremely full-featured help section that can bring you right up to speed about how backdoors work. We talk more about Port Explorer when we address sniffers in the next section — for now, firewalls are the order of the day; they can provide you information you can use to research suspicious port usage.

Tracking suspicious data flow

When you’ve established the presence of a suspicious outbound data flow, it’s helpful to examine its content and figure out where it’s going. That process has two vital stages:

- ✔ Use a sniffer to capture and inspect the individual data-transfer units (*packets*) to identify the information being transmitted.
- ✔ Try to map the Internet protocol (IP) address of the remote destination — to a specific locality and domain.

When you can get a close look at outbound packets, you can assess their legitimacy. Processes may open multiple sockets or ports on your computer to facilitate data transfer — legitimate or not. A *port-to-process mapping* program and a dedicated sniffing program are among the best tools for finding out what’s going on. These can

- ✔ **Resolve IPs to identify their domain names and countries of origin.** If you want to determine whether a domain is dodgy and/or associated with malware distribution, there are a number of methods you can use to research that: Google, SiteAdvisor, and an MVPS `HOST`-file search. If some countries turn up that have no legitimate connection with your business, that’s a big (and maybe sinister) hint of possible foul play.
- ✔ **Provide traffic-volume statistics.** Such figures reveal how many bytes and packets have been sent and received by specific ports on your network.



Chasing the dodging malware

Here's a real-life example of looking beyond the firewall. Nancy was recently assisting an advanced user whose computer was infected with a devious breed of malware — the threat was hopping from process to process, hoping to find one that would allow Internet access. It jumped from WinLogon to Explorer to `svchost` to Internet Explorer (which wasn't the default browser) — and later even his firewall and antivirus processes were used. It takes a sharp computer user to realize that there's something *wrong* with that pattern of activity, or even to be able to see it. Now, how many users would block Explorer or WinLogon or Internet Explorer — all established Microsoft processes — from gaining Internet access if their firewalls prompted them for approval? We'd bet most users wouldn't — but (luckily) the astute owner of this infected computer did.

The first clue that alerted him to this illegitimate activity was his firewall log — it revealed that all these valid Microsoft processes had made repeated attempts to connect to the same domain. (Hmmm . . .) After establishing that the domain was well known for distributing malware, he used a port-mapping program to monitor the illicit port activity — and a sniffer to trap its content. This example illustrates why it's important to assemble *all* the forensic evidence from various sources (firewall logs, port-mapping program logs, and so on) — in combination with a healthy dash of common sense — when you examine your system for malware clues. You have to piece the entire puzzle together to formulate the correct diagnosis — and it ain't easy! In the section called "Examining the firewall," we show you how you can use these same strategies to research suspect activity on your own system.

Catching rootkits phoning home

As discussed in Chapter 4, a firewall is the barrier between your computer or network and the World Wide Web. Most networks use both hardware and software firewalls. Routers come with two essential tools built in: *network-address translation* (NAT) capability and a firewall.



Although this dual layering provides extra protection, it doesn't make you resistant to attack. Many threats rely on clever social-engineering tricks to gain entry to a system (refer to Chapter 1) — not just vulnerabilities or improperly secured systems. When a pernicious threat has landed on your computer or network, you've got trouble. As we've learned, very crafty malware will milk any means available to connect remotely — and it often succeeds, even on a well-fortified system. Some types of malware inject a thread or DLL into a process that has legitimate access. Others use a reserved port

(SMTP, HTTP, DNS) that firewalls normally give the green light; others may just disable your security programs altogether to get a free pass (see the accompanying sidebar for a case in point).

Examining the firewall

One method of assessing intrusion is examination of firewall activity. Rootkits commonly hide a backdoor on an infected host. The backdoor establishes a remote connection to a remote client computer, from which it receives its directives. But outbound access attempts are normally blocked unless you've approved the process that's requesting an Internet connection. Of course, the firewall log contains a record of *all* blocked access attempts, to or from your system — so the firewall log can provide vital forensic clues. You can examine it to determine which processes have been trying to establish a remote connection to send out data packets from your computer or server to a remote destination.



To monitor outbound access attempts, you have to use a bidirectional firewall — *not* the Windows Internet-connection firewall, which provides only inbound protection. If your firewall isn't bidirectional, it can't fully monitor unauthorized programs residing on your system. Nor can it prevent them from *phoning home* — trying to establish an Internet connection so they can transmit data acquired from the host computer back to a remote server, reporting back to their faraway masters for purposes unknown.

This section concentrates on the symptoms that can signal rootkit behavior, and shows you how to troubleshoot those symptoms to determine whether they're indeed malicious. The practical place to start is to get a handle on how your firewall can provide valuable evidence as you troubleshoot suspicious process activity.

All software firewalls produce logs that record blocked attempts at inbound and outbound Internet access. Examining this log can give you insight into rogue processes that are trying to phone home. Programs that do this can vary in threat potential. Some may be adware- or spyware-oriented, in which case they're trying to report on your browsing habits so their controllers can deliver targeted advertising. Others, of a more insidious nature — backdoor trojans, for example — can send out highly confidential information and listen for remote commands that their far-off masters send to direct the trojans' execution. Regardless of the reason for this blatant invasion of privacy, you should make a habit of inspecting your firewall logs for any red flags that signal suspicious behavior.

Most firewall logs will identify the process that's calling home. This enables you to investigate the process to determine whether it's suspicious. However, as you can see in Figure 8-4, a process that's calling home may be hiding behind a legitimate `svchost` process. (As discussed in Chapter 7, `svchost.exe` is the Windows operating system's process that loads services very early in the boot sequence.)

If you inspect your active process list with Task Manager, it's normal to see several instances of `svchost.exe` running. However, since many forms of malware — particularly rootkits — install a malware service, it's possible that a blocked outbound `svchost.exe` entry in your firewall log may represent a malware service attempting to phone home. Identifying the guilty module or service that's running behind the `svchost` requires a little extra work.

The next several sections walk you through the general steps required to identify the process started by a `svchost.exe` that's listed as a blocked entry in your firewall. In our example, we use two programs: TCPView (a port-to-process-mapping program), and Process Explorer (an advanced process viewer and monitor) to assist our troubleshooting. Figure 8-4 illustrates a ZoneAlarm Firewall log that contains several blocked outbound access attempts.

The screenshot shows the ZoneAlarm Alerts & Logs window. The main area displays a table of firewall events. The table has the following columns: Type, Protocol, Program, Source IP, Destination IP, and Direction. The events listed are:

Type	Protocol	Program	Source IP	Destination IP	Direction
Firewall	UDP	svchost.exe	192.168.1.100...	167.206.251.4:53	Outgoing
Firewall	UDP		192.168.1.100...	192.168.1.102:138	Outgoing
Firewall	TCP (flags:S)		192.168.1.102...	192.168.1.100:139	Incoming
Firewall	TCP (flags:S)		192.168.1.102...	192.168.1.100:139	Incoming
Firewall	TCP (flags:S)		192.168.1.102...	192.168.1.100:139	Incoming
Firewall	UDP		192.168.1.100...	192.168.1.102:138	Outgoing
Firewall	UDP		192.168.1.100...	192.168.1.102:138	Outgoing
Firewall	TCP (flags:S)		192.168.1.102...	192.168.1.100:139	Incoming
Firewall	TCP (flags:S)		192.168.1.102...	192.168.1.100:139	Incoming
Firewall	UDP	svchost.exe	192.168.1.100...	167.206.251.4:53	Outgoing
Firewall	UDP		192.168.1.100...	192.168.1.102:138	Outgoing
Firewall	TCP (flags:S)		192.168.1.102...	192.168.1.100:139	Incoming

The 'Entry Detail' pane at the bottom shows the following information:

- Description: Packet sent from 192.168.1.100 (UDP Port 1138) to 167.206.251.4:53
- Rating: Medium
- Date / Time: 2006/11/02 12:29:18-5:00 GMT
- Type: Firewall

Figure 8-4:
The
ZoneAlarm
firewall log.

Steps for identifying a process

To trace which process is running behind the highlighted blocked `svchost.exe` entry listed in Figure 8-4, we ran through the series of steps outlined in this section.

Step 1: Identify the port

Right-clicking the first blocked `svchost.exe` entry in the ZoneAlarm firewall log allows you to copy the information for that entry to the Clipboard. A sample of those results follows:

Description	Packet sent from 192.168.1.100 (UDP Port 1138) to 167.206.251.4 (DNS) was blocked
Rating	Medium
Date / Time	2006/11/02 12:29:18-5:00 GMT
Type	Firewall
Protocol	UDP
Program	svchost.exe
Source IP	192.168.1.100:1138
Destination IP	167.206.251.4:53
Direction	Outgoing
Action Taken	Blocked
Count	1
Source DNS	ADMIN
Destination DNS	dhcpl.srv.whplny.cv.net

What is significant in this entry — and what allows us to trace the entry back to its originating process — is that the local port number 1138 is identified both in the `Description` field and the `Source IP` field (after the semicolon).

Step 2: Identify the process ID associated with the identified port

This step requires that we use a port-to-process mapping program. In this example we'll use Sysinternal's TCPView to identify the Process ID (PID) that's associated with Port 1138.

In TCPView's display, the *process ID* (PID) is the number to the right of the process, right after the semicolon. We can copy the information for the entry associated with Port 1138 (identified under the `Local Address` column, after the semicolon) by right-clicking the entry and choosing copy from the context menu. Doing so produces the following results that identify the PID as 984. Above the actual output we've included labels to help you decipher it:

Process	PID	Protocol	Local Address	Remote Address
svchost.exe:	984	UDP	127.0.0.1:1138	*:*


Step 3: Identify the process(es) loaded by the svchost.exe with PID 984

Process Explorer enables you to identify the process(es) or services loaded by any running `svchost` by merely hovering your cursor over the `svchost.exe` entry in the *process tree* (hierarchical listing of active processes). Since we know the `svchost.exe` we're interested in has a PID of 984, we can locate the process that corresponds to that PID in the process tree.



Any processes or services loaded by that `svchost` are identified in the tool tip. You can also view them by right-clicking the `svchost.exe` that matches the PID entry and then choosing Properties. By clicking the Services tab in the Properties dialog box that appears, you can identify the process(es) or service(s) loaded by this instance of `svchost.exe`. (Remember — a single `svchost.exe` can load one or more services.) Identification via this tool-tip method is depicted in Figure 8-5.

Figure 8-5:
Using
Process
Explorer to
identify
modules
loaded by
`svchost.exe`.



True, this example is a simulated exercise for the sake of illustration, and our `svchost.exe` in Process Explorer wasn't the real process that our original Zone Alarm `svchost.exe` really mapped to. But we were only unable to trace the identical `svchost.exe` entry because Process Explorer displays only *active* processes; the process that corresponded to the `svchost.exe` entry in the firewall log was no longer running. Therefore Process Explorer (in this case, anyway) wouldn't have been able to identify the entry. If, however, the entry had been an *application* process (as opposed to a service launched by a `svchost`), the process name would have been visible in your firewall log — with no further tracing required.



You have to catch the `svchost` process when it's still active in order to trace it to its underlying process or service when you're using Process Explorer. (For an example of a different real-time approach, see the nearby "Tracking down outbound access with Filemon" sidebar.)

Because our example was a simulated exercise on a clean computer, no disguised malware was revealed. But you can apply the same troubleshooting procedure to uncover malware processes and services on an infected system.



After you've identified the file on disk that matches the process you're researching, you can test it for its threat potential by uploading it to Virus Total or Jotti — two separate online services that employ over 20 different antivirus scanners to give you a cross-section of opinions in determining a sample's threat status. This resource is extremely useful for catching new threats; not all antivirus programs have the same definitions, nor do they update their databases at the same time. Here are the sites:

- ✓ Virus Total: www.virustotal.com/flash/index_en.html#01
- ✓ Jotti: <http://virusscan.jotti.org>

Tracking down outbound access with Filemon

Filemon is another Sysinternals tool that can help catch erratic processes in the act. It records a log of process activity in real time (as it's happening). To trace the suspect process in our example, however, Filemon would have to have been running at the moment the firewall alert occurred. Filemon is helpful in tracing the time sequence of malware installation — whether for testing purposes or for system troubleshooting. Filemon can play a crucial role in capturing and tracing malware activity, if you can identify some key event that triggers the erratic malware to become active. If that event is identified, then Filemon can be started up right before you initiate the triggering event, so the malware can be caught red-handed.

For example, if an outbound connection is established every time you play a particular game on your computer, then start Filemon, run the game, and trace whatever process is making that outbound connection. This informative exercise can determine whether the game came bundled with PUPs (potentially unwanted programs) or whether it always establishes a connection to support multiplayer gaming, even when that isn't required — good to know either way. You can use the filtering function to refine and focus Filemon's abundant output. Here's where to get more information on Filemon:

[www.sysinternals.com/
Utilities/Filemon.html](http://www.sysinternals.com/Utilities/Filemon.html)

Mapping the recipient IP address by doing a reverse DNS search

Tracing the destination IP address to its domain by doing a reverse DNS lookup (as detailed in the “Defending your ports” section earlier in this chapter) can provide valuable insight into whether the process that’s calling out is malicious. In our troubleshooting exercise, the IP was safe — we’ve often mapped IPs to dubious destinations. In most cases, if you have good reason to suspect the entity identified by the reverse DNS information, evidence to that effect will show up right away. So add this to your list of standard forensic procedures: mapping the destination IP (remote address) of a blocked outbound firewall entry to its domain.

There are several ways to assess the safety of a remote domain that shows up in your firewall log, and we outline those below. This will help you determine whether a remote connection attempt is malicious or not — and this will in turn help you establish the legitimacy of the process that is making the connection attempt.



Legit processes can be hijacked to acquire a remote connection (as was pointed out by our sidebar “Chasing the dodging Malware” real-life example, discussed earlier in this chapter) — this can complicate the picture when you’re trying to identify and remove malware that’s responsible for “phoning home.”

- ✓ Google the domain name. Sometimes just the Google results will yield enough information to determine whether the domain is safe.
- ✓ Visit McAfee SiteAdvisor’s homepage (www.siteadvisor.com) and plug the domain name into their *Look up a site report* service. SiteAdvisor returns immediate feedback on any Web site you enter (providing it’s in their database). SiteAdvisor is discussed in more detail in Chapter 4.
- ✓ See whether the domain is blocked by the MVPS host file. A blocking host file can be installed on any Windows computer; it automatically prevents its user(s) from navigating to known undesirable destinations. The MVPS HOSTS file contains an updated list of malicious sites in text-file format; you can open it in Notepad and search it for any suspicious domain name you choose to investigate (and if you get a match with one that’s listed in the MVPS HOSTS file, that can’t be good!). You can download, view and search the MVPS HOSTS file at the MVPS.org Web site (www.mvps.org/winhelp2002/hosts.txt).

Researching process databases

The troubleshooting example just provided illustrates how you can trace attempts to phone home back to their originating processes. If the process

were still active, we would Google it or search a process database to determine if it were harmful or not. Most likely, a malware or rootkit-related process would still be running until you stopped it (and you might even have trouble doing that — those processes can be very stubborn). The following process databases can help you research and identify processes:

- ✓ Answers that work: — Task List programs (includes services)

www.answersthatwork.com

- ✓ The Process Library:

www.processlibrary.com

- ✓ WinTasks Process Library

[www.liutilities.com/products/wintaskspro/
processlibrary](http://www.liutilities.com/products/wintaskspro/processlibrary)

- ✓ CastleCops Services List (this list identifies stealth malware services with a red warning label)

<http://castlecops.com/O23.html>

Windows Startup List Databases (that list programs and services that can be set to automatically run at boot up) can also provide important identifying information:

- ✓ Bleeping Computer Startup Programs Database:

<http://www.bleepingcomputer.com/startups/>

- ✓ CastleCops StartupList (affixes a red rootkit type stealth involved note to identify rootkit startups):

<http://castlecops.com/StartupList.html>

- ✓ Gratis Application Database:

<http://www.greatis.com/appdata/>

Trusting Sniffers and Firewalls to See What Windows Can't

When information is sent across the Internet, it's transmitted in small units called *packets*. Most communications (such as e-mails) must be broken up into a series of packets — which are then reassembled into their original

form when they reach their final destinations. The use of small individual packets helps maintain and control adequate data flow across networks. *Sniffers* are programs used to capture and inspect the packets that flow between networked computers. The following sections describe how sniffing works — and how you can use it against blackhat hackers.

How hackers use sniffers

It's not uncommon for an intruder to install a sniffer after gaining first entry to a system. An attacker uses a sniffer to collect all the information that passes through a computer's local network segment. The collected information then goes through sifting and analysis to find something valuable to the intruder — such as usernames, passwords, account numbers, and other confidential information. An especially coveted prize would be an *administrative* password, which would grant root access to the entire system — along with all administrative privileges (in effect, free rein to do anything on the network). Any useful information harvested from the sniffed data is usually inserted into a file that's then transmitted back to the hacker via a backdoor. Sniffers can create huge logs on the hard drive of the infected system containing all the information they have sniffed. The data captured in the logs is inspected for any information that might prove fruitful. These logs are usually massive — and if you can find them, they can tip you off to a successful break-in. Bottom line: If there is an inexplicable increase in demand for hard-drive space, it's a condition that merits investigation. Sooner rather than later.

Sniffers can sample all the information passing to and from a networked computer by manipulating the Network Interface Card (NIC) on a compromised system so it goes into promiscuous mode. This enables the sniffer to collect the information from *all* data packets before they're delivered to their intended destinations. This sniffing process is passive — the data is sampled while in transit, without being altered in any way — which makes the process difficult to detect.

Using sniffers to catch hackers at their own game

In a strange double twist, a sniffer can be used in two opposing ways, by people with mutually exclusive goals:

- ✔ **To exploit a network:** hackers can use sniffers to intercept and read privileged communications across a local network.
- ✔ **To catch a hacker:** System administrators can use sniffers legitimately to examine the content of data that flows out from a network, looking for signs of malicious intent.

You can use sniffers to capture data that passes through any network interface. Each data packet contains a *header* — a place to put the information that ensures the packet is delivered and recombined properly. The header information contains the source and destination IP addresses, as well as the source and destination port numbers. It works like the sender and recipient addresses you see on an envelope that contains a letter delivered through the post office. A packet's header information can provide valuable clues for assessing whether the packet should be considered suspicious. Then the content of the packet can be analyzed to see what's being sent.



The objective of an intruder is to retrieve readable, plain-text information of value from the local network. Packets containing encrypted data are not readable to anyone who lacks the decryption key. This makes encryption one of the best defenses a network administrator can employ against a blackhat's attempt to sniff the network. Therefore using transfer protocols that support password encryption is very important.



The FTP, IMAP, NNTP, POP3, SMTP-AUTH, and Telnet transfer protocols *do not* support password encryption; they're vulnerable to exploitation.

Normally, packet sniffers are a useful defensive tool if some indication of suspicious activity crops up on a network (or even a home PC). For example, if a program on the host computer is engaged in downloading material, you could use a sniffer to assess whether the incoming material is suspicious. You could also use it to determine the source IP address of the download — and then you could trace it to its domain name and physical location by performing a reverse DNS lookup. Alternatively, if you have a port-to-process mapping program and it indicates activity on a port that belongs to an unknown or questionable process, you could enable sniffing on that specific port or process — and then analyze the collected data to assess whether the transfer is legitimate.

Testing to see whether your NIC is in promiscuous mode

One method of detecting illicit sniffing is by testing to find out whether any adapter cards on a network are in promiscuous mode. An interface that's in

promiscuous mode is listening to *all* network traffic — which usually signifies that a network sniffer is actively installed.

Several available tools can detect network adapter(s) configured to run in promiscuous mode. Here are two such tools that can be used on Windows 2000 and XP systems:

- ✓ **PromiscDetect.** This is a freeware command-line utility that can be downloaded here:

```
http://ntsecurity.nu/toolbox/promiscdetect/
```

- ✓ **AntiSniff.** This popular sniffer-detection program is available for a free trial download at this address:

```
www.securitysoftwaretech.com/antisniff/
```



If you run a sniffer-detection program and get negative results, that doesn't prove conclusively that a sniffer does *not* exist. As you know, malware — particularly rootkits — can intercept and modify the results returned by any utility. If the tool does detect the presence of a sniffer — and no other factor seems to be causing those results — then you have definite reason for concern. But sometimes the tools themselves can create false positives; VMWare, for example, puts NICs in promiscuous mode as part of its configuration setup.

Sniffers you can use

Two programs you can use to sniff out hackers are the venerated and multi-talented Port Explorer and Ethereal. The following sections give you a closer look at each.

Port Explorer

Port Explorer was discussed previously for its superior port-to-process mapping capabilities. Port Explorer can also be launched from the run line through AntiHookExec to reveal hidden user-mode rootkit ports and the processes associated with them.

We like to use Port Explorer in combination with Process Explorer to determine the processes that underlie any `svchost` entries. `Svchost` entries in the Port Explorer display are identified as “Generic Host Process for Win32 Services” under the Process column. By using the Process ID (PID) in the next column,

we can easily see what services have been launched by this `svchost` by matching this PID to the `svchost` with the same PID in Process Explorer.

Process Explorer has the option of performing signature verification on any process or driver. When the `System` process is selected in the Process Explorer Process tree and the signature verification option is enabled, Process Explorer will list all kernel drivers along with their verification status in the lower pane. Process Explorer uses the Internet to obtain this digital signature data. This action results in a burst of port activity — which Process Explorer initiates to obtain timely verification. Figure 8-6 illustrates the Port Explorer main display, revealing all the open ports Process Explorer has opened to verify driver signatures.

Port Explorer has a very convenient Spy feature that allows you to selectively capture packet data on an individual port or entire process. When you select a process or port listed in the display and then right-click the entry, a context menu with a list of options appears. By choosing Process ⇨ Enable Spying from the context menu, we can capture packet data on the selected process. Alternatively, we could have chosen Socket ⇨ Enable Spying to capture data on a single port. We've elected to enable spying on `procexp.exe` (the Process Explorer program's executable file) as it verifies driver signatures. This enables us to collect packet data on every port Process Explorer has opened. We did this solely to illustrate Port Explorer's sniffing capabilities in the absence of having any trojan process to spy on. If we'd had an active backdoor trojan process, we'd have chosen to spy on that for sure.

Process	ProcessID	Country	Protocol	Local Address	Local Port	Remote Address
TrueVector Service	296		IPv6	localhost	0	192.168.1.1
LSA Shell (Export Version)	648		UDP	localhost	500	*.*.*
LSA Shell (Export Version)	648		UDP	localhost	4500	*.*.*
WebProxy	2768		UDP	localhost	18001	*.*.*
WebProxy	2768		UDP	localhost	18002	*.*.*
WebProxy	2768		TCP	localhost	31595	localhost
Generic Host Process for Win32 Services	3168		UDP	SOPHIA	123	*.*.*
Generic Host Process for Win32 Services	3168		UDP	localhost	123	*.*.*
Generic Host Process for Win32 Services	3168		UDP	localhost	4449	*.*.*
Generic Host Process for Win32 Services	3168		UDP	localhost	4450	*.*.*
Generic Host Process for Win32 Services	3168		UDP	localhost	4454	*.*.*
Generic Host Process for Win32 Services	3168		UDP	localhost	4455	*.*.*
Generic Host Process for Win32 Services	3168		UDP	localhost	4584	*.*.*
Generic Host Process for Win32 Services	3168		UDP	localhost	4585	*.*.*
Sysinternals Process Explorer	4088	United States	TCP	SOPHIA	4855	crl.microsoft...
Sysinternals Process Explorer	4088	United States	TCP	SOPHIA	4856	crl.microsoft...
Sysinternals Process Explorer	4088	United States	TCP	SOPHIA	4859	crl.microsoft...
Sysinternals Process Explorer	4088	United States	TCP	SOPHIA	4860	crl.microsoft...
Sysinternals Process Explorer	4088	United States	TCP	SOPHIA	4865	crl.microsoft...
Sysinternals Process Explorer	4088	United States	TCP	SOPHIA	4866	crl.microsoft...
Sysinternals Process Explorer	4088	United States	TCP	SOPHIA	4867	crl.microsoft...
Sysinternals Process Explorer	4088	United States	TCP	SOPHIA	4868	crl.microsoft...

#	Time	Process:PID	Action	Protocol	Local...	Remote Address
20042	15:50:49	WebProxy.exe:2768	RECEIVE	UDP	local...	localhost:0

Figure 8-6:
The Port Explorer main display.

Figure 8-7 illustrates the contents of a data packet captured using Port Explorer: The packet reveals a large amount of information that can be very helpful in tracing the nature of the data exchange — and can help establish the purpose of an active process with an open port. For example, the packet shows that Process Explorer connected to a Microsoft server in Redmond, Washington to obtain signature-verification data. More specifics are listed as well.

If we had enabled spying on a true trojan or a suspect process, we could have used Port Explorer's Disable Sending and Disable Receiving options to suspend data transfer to and from the process. Although this would block the process's dangerous data-transfer capabilities, it would not prevent spying on the process.




Figure 8-7:
Port
Explorer
driver
verification.

If analysis of the captured packet data determined that the process was indeed malicious, we could then kill it by selecting it and choosing Process → Kill Process from the context menu.

Port Explorer also offers Whois mapping, pinging, tracing, and resolving of a remote IP address to its domain and geographical location. You can access

these functions by choosing them from the context menu or by highlighting a displayed process and then choosing Options on the Menu bar. You can update Port Explorer's port-and-domain database by choosing Help ⇨ Check for ⇨ New Port and Domain Databases.

Ethereal

Port Explorer is a port-to-process mapping program with sniffing capability; Ethereal is a dedicated sniffer. It's a highly recommended freeware program that's very effective at capturing and analyzing packet data on networks. Ethereal can be downloaded at www.ethereal.com.

Ethereal allows you to select the network interface you want focus on. You can then capture all inbound and outbound packets transmitted through the selected network interface. If a piece of malware is downloading something from a Web site, the download can consist of 100 or more different packets, each with its own entry in the Ethereal main display; these packets can be written to a capture file on disk. When you stop recording this data, Ethereal displays a list all the packets it has captured; each packet is numbered and appears on a separate line in the display, along with source address and destination IP address. You can then have Ethereal display the content and header of any packet by right-clicking the packet and selecting viewing options from the context menu.

Figure 8-8 shows Ethereal's main display after a number of packets have been collected; each numbered line represents one packet of data.




Figure 8-8:
Ethereal's
main
interface.

Investigating Lockups and Other Odd Behavior

Windows comes with utilities that can help you investigate your system when it seems to be behaving oddly. Two standard Windows utilities are very helpful in troubleshooting odd behavior — and even though they can't spot a hidden rootkit they may provide clues to its existence: Event Viewer and Task Manager. (A little farther along, we review some third-party tools that extend the basic features of these programs. We show you how to turn on event auditing in Chapter 6.)

As mentioned in Chapter 6, Windows keeps a history of the significant actions that take place on your computer and records each action in one of three different logs — Application, Security, or System, depending on the action that triggered the event:

- ✔ **Application:** Contains events specific to the software programs on your system — such as program-installation failures, program starts and stops, and application-update failures.
- ✔ **Security:** Contains events which can have an impact on the security of your system — for example, file-system changes, service installations, and logon events.
- ✔ **System:** Contains events relating to system activity that often happens in the background but can trigger on-screen alerts. Typical entries include time synchronization to update your clock, automatic update installation, or inability to renew IP-address notifications.

The Application and System Event logs classify event-log entries as belonging to one of three types — Warning, Error, or Information; the Security Event Log characterizes its entries as Successes or Failures only.

The Microsoft service that records these events is called (wait for it) Event Log, and it's viewable in the Services Console. Use the Windows Event Viewer utility to view the content of any of these three event logs.

Accessing Event Viewer

The best way to get the flavor of what each log contains is to open the Event Viewer and inspect the events listed there. You can access the Windows Event Viewer in one of the following ways:



- ✓ If your system's Control Panel is configured to the Classic View, you access Event Viewer by choosing Start ⇨ Control Panel, double-clicking Administrative Tools, and then double-clicking Event Viewer.
- ✓ If your system's Control Panel is configured to the Category View, you access the Event Viewer by choosing Start ⇨ Control Panel⇨ Performance and Maintenance ⇨ Administrative Tools, and then double-clicking Event Viewer.
- ✓ If you want to access Event Viewer quickly from the run line, choose Start ⇨ Run, typing `eventvwr.msc` into the Open: field, and pressing Enter (or clicking OK).

Making some necessary tweaks to streamline logging

Right out of the box, Windows doesn't allow a very large default size for its event logs, nor does it *autoarchive* your event logs (automatically back them up so their data won't be lost). The Windows XP and Windows 2003 Server platforms allocate different amounts of space for event logs because of their different usage expectations (networks are busier than standalone computers). Event Viewer has a function that allows you to adjust the size of any event log to fit your requirements; we show you how to modify that setting both manually and with a small program we've created (to expand the maximum size beyond what you can get through manual adjustment alone).

Event logs are overwritten once their maximum size is reached (they're filled to capacity). This won't hold you in very good stead if you're in the middle of a forensic investigation — tracking down clues — and the event data suddenly becomes nonexistent! Setting the event logs to autoarchive is the remedy for that. The following sections show you how to reconfigure Event Log settings — both manually (in Event Viewer) and by using a Registry script to exceed the maximum limits available in Event Viewer. (Microsoft has documented the latter method, so don't worry — it's legit.)



These modifications are more appropriate for network servers that support multiple client computers — or for busy smaller networks — than for standalone computers. If you're thinking of modifying a single computer's settings, manually adjusting the maximum size in Event Viewer (or even leaving the Windows default setting as is) may be perfectly adequate.

Changing the default log size

In this section, we show you how to change the default Event Log size by giving you suggested settings that you can customize to fit the needs of your

own computer setup. Keep in mind, the example we give here is only for illustration purposes; your settings will be specific to your system.

1. **Choose Start ⇨ Control Panel.**
2. **Open Event Viewer, using the method that matches the view you've chosen for Control Panel.**
 - **Classic View:** Double-click Administrative Tools and then double-click Event Viewer.
 - **Category View:** Choose Performance and Maintenance ⇨ Administrative Tools, and then double-click Event Viewer.

Event Viewer appears.

3. **Right-click Security and choose Properties.**
4. **In the Log size box, change Maximum log size to 51200.**

This value sets the size at 50 megabytes, which allows 150,000+ entries in the log.

In Windows 2000 and Windows XP, the Event Logs default setting for Maximum Size is 512KB. In Windows Server 2003, the default for Maximum Size is 16MB, but these can be adjusted as we've shown you in our example.



Depending on your computer's configuration (business, home user, network), event logs can fill up quickly. Under certain normal circumstances — application installations, file creation and deletion, service starts and stops, running scheduled tasks, busy logon activity, software updating — logging activity can intensify. When an event log's maximum capacity is reached, the event-log data is overwritten. Normal system activity alone produces many event log entries, but if “failure” events occur, then your event logs will incur even more than the usual number of entries — because either the OS or the application (depending on the source of the failure) automatically attempts to rectify the error, and those attempts are also recorded in the event log. In a home or small-office situation, 150,000 entries should be adequate for at least 10 days' worth of log entries. If you see that your log is overwriting the oldest entry daily (or sooner), then it's time to reevaluate and reconfigure your Event Log settings.

Automatically archiving logs

Normally events logs are overwritten if they're more than a week old, but there is a way to make sure event logs are automatically archived to prevent you from losing data that can be vital to an investigation.



The event logs on large networks fill up very rapidly; all their data is lost when they reach their maximum size and are overwritten. To prevent Event Log data loss, we've got a Registry workaround: that sets the event logs to autoarchive so their content is preserved and resets the values for `Retention` and `MaxSize` (don't worry — we've created a script to do all that for you). After the script runs, the Event Logs autoarchive every time the log file reaches 100 MB. The Registry script that accomplishes this handy bit of hocus-pocus is posted on the CastleCops Rootkit Revelations Forum, located at this link:

```
www.castlecops.com/f233-Rootkit_Revelations.html
```

We've posted the script's content here to give you a preview of exactly what it does before you download and run it:

```
Windows Registry Editor Version 5.00


[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\EventLog\Application]
"MaxSize"=dword:06400000
"Retention"=dword:00278d00
"RestrictGuestAccess"="1"
"AutoBackupLogFiles"=dword:00000001
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\EventLog\Security]
"MaxSize"=dword:06400000
"Retention"=dword:ffffffff
"RestrictGuestAccess"="1"
"AutoBackupLogFiles"=dword:00000001
"WarningLevel"=dword:0000005a
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\EventLog\System]
"MaxSize"=dword:06400000
"Retention"=dword:00278d00
"RestrictGuestAccess"="1"
"AutoBackupLogFiles"=dword:00000001
```

Because Event Log is a Microsoft service responsible for recording your event logs, it appears in the Services Console (as shown in Figure 8-9).

The Event Log service uses these Registry keys:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
    Eventlog\Security
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
    Eventlog\Application
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
    Eventlog\System
```

Figure 8-9:
The Event
Log
Services
Console.



Although the `MaxSize` and `Retention` subkeys exist for each Event Log service key just listed, the `AutoBackupLogFiles` does not — it must be created by using the Registry script. If you look at the following list, you'll see what happens to the Security Event Log Registry keys after you run the Registry script; the `AutoBackupLogFiles` subkey appears in the third entry:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
    Eventlog\Security\Retention
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
    Eventlog\Security\MaxSize
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
    Eventlog\Security\AutoBackupLogFiles
```

Inspecting event logs with Windows Event Viewer

Chapter 6 describes the Windows Event Log service and shows you how to turn it on; in this section, we show you how to use Event Viewer to examine the information that the event logs provide when you're tracking down rootkit activity or evidence of intrusion.



In Windows XP Professional and Windows 2003 Server, the auditing of Security events isn't turned on by default; to use it (if your operating system is one of those versions) you have to turn it on and configure it, using the settings in Chapter 6 as a guideline.



Because rootkits conceal their presence by hooking operating-system APIs, you have to consider that they can hook event-logging APIs as well — and erase the evidence of their existence. Rootkits are known to intercept and alter the following Event Log APIs to falsify Event Log reports:

```
EventLog.GetEventLogs  
EventLog.WriteEntry  
EventLog.Delete
```

Even if this happens, however, the log entries created from original break-in should remain intact — and they can supply valuable evidence to substantiate the attack.

You can research Event Log entries to trace events that caused noteworthy system changes. If you suspect an intrusion has taken place, you can use the entries to obtain troubleshooting information so you can reconstruct a specific incident or weave a trail of evidence intrusion.

It's good security practice to periodically inspect your event logs to monitor system activity for any anomalies (which may not be readily apparent from your computer console). This can help you hone in on potential problems before they have escalated beyond your control. For example, system administrators may want to routinely audit failed logon attempts due to incorrect usernames or passwords. Such events might suggest there's an intruder trying to access a system by guessing passwords.

The Event Viewer enables you to inspect all three types of event logs — Application, Security, and System. These logs are in hidden files located in the %Windir%\system32\config folder and these are their paths (if your OS is located on drive C:):

```
C:\WINDOWS\system32\config\SECURITY.LOG  
C:\WINDOWS\system32\config\SOFTWARE.LOG  
C:\WINDOWS\system32\config\SYSTEM.LOG
```

Figure 8-10 shows how these same files appear in Windows Explorer.

When you open the Event Viewer you will see a window like that shown in Figure 8-11.




Figure 8-10:
Event Log
locations.




Figure 8-11:
Event
Viewer.

The left pane lists the Application, Security, and System event logs in hierarchical fashion:

- ✓ The Application Event log shows application installations, application information alerts, critical errors, and auditing errors.
- ✓ The Security Event log shows only successful and failed audits, based on the nine events you've elected to audit in your security audit configuration (as discussed in Chapter 6).

- ✓ The System Event log shows services being started and stopped, stop errors, device errors, Windows Update installations, network communication errors, and system warnings such as `Windows is increasing the size of your Virtual memory`. Three classifications of system events are available: Information, Warning, or Error.

Clicking any individual event-log type displays the contents of that log in the right pane: Double-clicking an event entry (or right-clicking it and selecting Properties) brings up a Properties dialog box with more information about the event that triggered the entry.



By default, the display is in time-sequence order, with the most recent event appearing first on the list. This can be changed by clicking the column header you choose, to sort the events.

Investigating Event Viewer results

It's a clear security priority to investigate noteworthy events that suggest intrusion or evidence of rootkit behavior. The first practical step here is to inspect the logs and filter their results to show EventIDs that may be telltale signs of malware.

Some events that can indicate system intrusion or malware penetration include the following:

- ✓ **Unsuccessful logon attempts.** Table 8-1 shows the EventIDs that might signal intrusion attempts.

Table 8-1	Event Log IDs Suggestive of Intrusion
<i>Account-related IDs</i>	
624	Create User Account
627	Change Password attempt
628	User Account / Password reset
630	Deleting a User
642	Changing a User Account
685	Changing an Account Name
<i>Account Lockouts</i>	
529	Logon failure due to incorrect username or password

(continued)

Table 8-1 (continued)

<i>Account Lockouts</i>	
644	A user account was autolocked
675	Pre-authentication failed on a DC (incorrect password)
676	Authentication ticket request failed
681	Logon failure

- ✔ **New installations or service starts.** Because most rootkits employ a kernel-level driver that's installed as a service, searching the System Event Log for new service installations or starts can help you assess a key indicator of possible rootkit infection. Another indicator would be new hardware-driver installations.
- ✔ **Stop errors.** Drivers installed by third-party software can cause stop errors — an error that produces a Blue Screen of Death (BSOD) — which in turn can be symptomatic of rootkit infection. Rootkit drivers must interact intricately with the host operating system's kernel. If they're coded improperly, system instability and crashes can occur. When this happens, you get a BSOD, a fatal-error message, an error code — and a reason for suspicion.

All stop errors are listed in the System Event Log. Stop-error codes generated by faulty rootkit drivers usually have a specific stop-error code: 0x00000050. Normally you also get a System Event Log entry with `Error EventID = 1003`. Figure 8-12 shows an example of an Event Log containing a stop error.



MonitorWare maintains a list of the Top 50 Events that trigger Event Log entries. Here's where to find the list:

www.monitorware.com/en/events/top50events.php

Filtering log data

Filtering is a useful function that's accessed by choosing View ⇨ Filter in the Event Viewer. Filtering allows you to narrow down your results by applying selection criteria. You can filter events in several useful ways:

- ✔ **By user:** List only those events created by a specific user.
- ✔ **By event type:** For example, you can filter for Information, Warning, Error, Success Audit, and Failure Audit events.
- ✔ **By specific application or system service**




Figure 8-12:
A stop
error in
the System
Event Log.

You can also enter an EventID value to select entries that correspond to a specific EventID. (Figure 8-13 shows the dialog box for choosing your filtering options.)




Figure 8-13:
The Event
Logs Filter
dialog box.

Now let's see how filtering might be applied to select the results according to a specific EventID. Figure 8-14 shows an example of a Security Event Log that looks fairly typical, but notice that an "anonymous" user has successfully logged on — which corresponds to EventID 540.

Figure 8-14:
The Security
Event Log
showing an
anonymous
logon
attempt.

Type	Date	Time	Source	Category	Event	User
Success Audit	6/14/2006	1:21:26 AM	Security	Logon/Logoff	528	LOCAL SERVICE
Success Audit	6/14/2006	1:20:07 AM	Security	Privilege Use	576	NETWORK SERVICE
Success Audit	6/14/2006	1:20:07 AM	Security	Logon/Logoff	528	NETWORK SERVICE
Success Audit	6/14/2006	1:19:18 AM	Security	System Event	515	SYSTEM
Success Audit	6/14/2006	1:19:05 AM	Security	Policy Change	615	NETWORK SERVICE
Success Audit	6/14/2006	1:19:04 AM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	6/14/2006	1:18:57 AM	Security	Privilege Use	576	CHECKLETTE
Success Audit	6/14/2006	1:18:57 AM	Security	Logon/Logoff	528	CHECKLETTE
Success Audit	6/14/2006	1:18:57 AM	Security	Account Logon	680	SYSTEM
Success Audit	6/14/2006	1:18:57 AM	Security	Logon/Logoff	538	CHECKLETTE
Success Audit	6/14/2006	1:18:57 AM	Security	Privilege Use	576	CHECKLETTE
Success Audit	6/14/2006	1:18:57 AM	Security	Logon/Logoff	528	CHECKLETTE
Success Audit	6/14/2006	1:18:57 AM	Security	Account Logon	680	SYSTEM
Success Audit	6/14/2006	1:18:57 AM	Security	Policy Change	806	SYSTEM
Success Audit	6/14/2006	1:18:55 AM	Security	System Event	515	SYSTEM
Success Audit	6/14/2006	1:18:55 AM	Security	System Event	515	SYSTEM
Success Audit	6/14/2006	1:18:53 AM	Security	System Event	515	SYSTEM
Success Audit	6/14/2006	1:18:52 AM	Security	Privilege Use	576	NETWORK SERVICE
Success Audit	6/14/2006	1:18:52 AM	Security	Logon/Logoff	528	NETWORK SERVICE
Success Audit	6/14/2006	1:18:52 AM	Security	System Event	515	SYSTEM
Success Audit	6/14/2006	1:18:52 AM	Security	System Event	518	SYSTEM
Success Audit	6/14/2006	1:18:52 AM	Security	System Event	515	SYSTEM
Success Audit	6/14/2006	1:18:52 AM	Security	System Event	515	SYSTEM
Success Audit	6/14/2006	1:18:52 AM	Security	System Event	515	SYSTEM
Success Audit	6/14/2006	1:18:52 AM	Security	System Event	515	SYSTEM
Success Audit	6/14/2006	1:18:52 AM	Security	System Event	514	SYSTEM

There is only one such attempt in the entire display that we can see — but suppose we want to see how many *other* successful anonymous logon attempts have happened. We could search for only those entries that match the 540 EventID by choosing View ⇨ Filter and plugging **540** into the EventID box. Figure 8-15 illustrates the successful anonymous-logon attempts.

Figure 8-15:
The Security
Event
Log with
Anonymous
Logon Filter
applied.

Type	Date	Time	Source	Category	Event	User
Success Audit	6/14/2006	1:19:04 AM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	6/12/2006	2:04:27 AM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	6/1/2006	8:47:09 PM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	6/1/2006	11:18:13 ...	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	5/22/2006	1:20:28 PM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	5/21/2006	1:12:44 PM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	5/13/2006	8:00:20 PM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	5/12/2006	1:16:35 AM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	5/7/2006	1:26:55 PM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	5/7/2006	10:27:42 ...	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	4/26/2006	4:32:38 PM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	4/22/2006	6:48:59 PM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	4/12/2006	2:38:29 PM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	4/10/2006	11:10:11 ...	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	3/27/2006	11:41:27 ...	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	3/26/2006	11:20:15 ...	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	3/19/2006	12:36:32 ...	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	3/18/2006	7:20:19 PM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	3/17/2006	4:09:57 PM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	3/15/2006	2:05:33 AM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	3/6/2006	4:43:32 PM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	2/26/2006	4:13:37 PM	Security	Logon/Logoff	540	ANONYMOUS LOGON
Success Audit	2/26/2006	3:37:41 PM	Security	Logon/Logoff	540	ANONYMOUS LOGON

Upgrading to Event Log Explorer

Here's where we demonstrate researching this entry further with a third-party alternative to the Windows Event Viewer called the Event Log Explorer. Event Log Explorer has the same features as the Windows Event Viewer — but extends its usefulness with a few add-ons. Event Log Explorer is a small freeware program, available for download at the developer's Web site (www.eventlogxp.com).

First, we bring up the same event log (we just viewed in the Windows Event Viewer), but this time, we'll view it in Event Log Explorer — and click the identical anonymous entry. Event Log Explorer allows you to copy and paste a log entry to the Clipboard, so it can be entered into reports, a text file, or a word-processing document: Here is the Clipboard information that Event Log Explorer generates for the anonymous entry:

```
Type: Audit Success
Date: 6/14/2006
Time: 1:19:04 AM
Event: 540
Source: Security
Category: Logon/Logoff
User: NT AUTHORITY\ANONYMOUS LOGON
Computer: SOPHIA
Description:
Successful Network Logon:
User Name:
Domain:
Logon ID: (0x0,0x14355)
Logon Type: 3
Logon Process: NtLmSsp
Authentication Package: NTLM
Workstation Name:
Logon GUID: {00000000-0000-0000-0000-000000000000}
```

Some of the nicer features of Event Log Explorer include these:

- ✔ You can copy event-log entries to the Clipboard, print entire event logs, and export them — in text, HTML, or Excel format. To copy a highlighted log entry to the Clipboard in Event Log Explorer, simply choose Event ⇄ Copy to Clipboard.
- ✔ Event Log Explorer extracts the description of the event that triggered the Event Log entry and lists it at the bottom of the page. (Windows Event Viewer requires you to double-click the entry and bring up a Properties dialog box before you can see this information.)

- ✔ To gather more information about any EventID, you can choose the “Lookup in Knowledge Bases” option from the context menu of any individual event-log entry. When you choose Event ⇨ Lookup in Event.ID.Net Database, you’re taken to an extensive EventID database compiled and maintained by Altair Technologies — with thousands of EventIDs from hundreds of event sources, thousands of comments and contributors. This vast resource can help you decipher and assess the legitimacy of an entry. Figure 8-16 shows the report generated when we plugged EventID 540 into the Event.ID.Net Database.
- ✔ You can also build a network tree to view the event logs on specific client computers attached to your network. To bring up a log from any of them, you need only choose it from the tree — and you can have multiple logs open at the same time.
- ✔ Event Log Explorer has a built in log archive feature; you choose File ⇨ Save Log As to save the current Event Log data so it won’t be lost when Windows overwrites the Event log files.
- ✔ You can selectively print event logs by choosing File ⇨ Print Options, and you can save them to a file, in the format you specify, by choosing File ⇨ Export.

Event ID: 540	
Source	Security
Type	Success Audit
Description	Successful Network Logon User Name: <user name> Domain: <domain name> Logon ID: <logon id> Logon Type: <logon type> Logon Process: <logon process> Authentication Package: <authentication package> Workstation Name: <computer name>
Things to understand	What is an authentication protocol?
Comments	<p>Adrian Grigora (Last update 5/20/2004): This event indicates that a remote user has successfully connected from the network to a local resource on the server, generating a token for the network user. For example, mapping a drive to a network share or logging with an account whose profile has a drive mapping would generate this auditing message.</p> <p>See the links to Windows Logon Types, Windows Authentication Packages and Windows Logon Processes for information about these fields. Understanding how the logon took place (through what channels) is quite important in understanding this event.</p> <p>This event may also be reported for builtin accounts. Whenever a user logs in the associated builtin accounts are also logged in. The HelpAssistant account in Windows XP is one such account. Even if the Remote Assistance Service is disabled, the account will still login. This is not a potential security violation as the HelpAssistant account itself is disabled. See M300692.</p> <p>lonut Marin (Last update 9/2/2005): This event informs you that a logon session was created for the user. For information on the details accompanying the event (logon ID, logon GUID, etc.) see MSW2KDE.</p> <p>See M287537, M326985, for additional information on this event.</p>
Links	M174074 , M287537 , M300692 , M326985 , Windows Logon Processes , Windows Logon Types , Windows Authentication Packages , Online Analysis of Security Event Log , MSW2KDE
Search	Microsoft Support - Microsoft Search - Google Groups - Google Microsoft - EventID.Net Query - More links...
Various	Send comments - Notify me when updated

Figure 8-16:
You can do a lookup like this in Event.ID.Net Database.

- ✓ The filtering feature lets you do searches on multiple EventIDs by specifying ranges — or by simply listing different EventIDs separated by commas. (Windows Event Viewer only allows a single EventID per search.) Searching on multiple EventIDs is helpful when a specific type of suspicious activity (for example, repeated logon attempts) can be attributed to more than one EventID type.

Trying MonitorWare

If you're still concerned about the entry and haven't gotten to the bottom of its cause, try another EventID database provided by MonitorWare:

`www.monitorware.com/en/events/index.php`

MonitorWare generates the report shown in Figure 8-17 when we plug in EventID 540 — indicating that the entry is a normal encounter, and not really anonymous but null. (A MonitorWare message characterizes the event as a typical occurrence: "It is normal to see these anonymous logins — they do not indicate somebody broke in.") That's a satisfactory result; the event is no longer suspicious — or, at this point, worth pursuing.

The screenshot shows a browser window titled "Security - Security - ID 540: Event Reference - MonitorWare - Opera". The address bar shows the URL "www.monitorware.com/en/events/details.php?details_id=14". The main content area displays a report for Event ID 540, structured as follows:

Event Type	Windows Event Log
Log Part	Security
Source	Security
Event ID	540
Meaning	Successful Network Logon. This event indicates that a remote user has successfully connected from the network to local resource on the server, generating a token for the network user.
Description	This indicates a successful logon. Please note that sometimes the user ANONYMOUS is the logged on user. Anonymous logon means that it is a NULL session. NT Auth/Anonymous is just a pseudonym for a NULL Session. User connections should never come in under NT Auth/Anonymous since this isn't really an account; it just means that no credentials were supplied. So it is normal to see these Anonymous logins — they do not indicate that somebody broke in.
More Informations	<p>Cause</p> <p>A logon session was created for the user. The message contains the Logon ID, a number that is generated when a user logs on to a computer. The Logon ID that is assigned to a logon session is unique to that logon session until the computer is restarted, at which point the Logon ID may be reused. The Logon ID can be used to correlate a logon message with other messages, such as object access messages.</p> <p>Resolution</p> <p>No user action is required.</p>

On the right side of the report, there are two advertisements: "Game Server Monitor" and "Event Log Monitor".

Figure 8-17:
A typical
Monitor
Ware
report.

Log Parser

Windows Event Viewer and Event Log Explorer are both handy for viewing and researching Event Log entries for single computers and smaller networks. If you're on a larger network, however, wading through *thousands* of daily Event Log entries to isolate suspicious activity can be a daunting prospect. That's where the Microsoft command-line utility called Log Parser can be an invaluable aid. Log Parser enables you to query the Event Log database (or any other database) for entries that match your specific criteria, using SQL scripts that use this command syntax:

```
LogParser <command>
```

Here is an example of a command that searches for successful logons — in this case EventID = 528 for a particular user ('%TESTUSER%):

```
LogParser "SELECT TimeGenerated, SourceName, EventCategoryName, Message INTO
report.txt From security.evt WHERE EventID = 528 AND SID LIKE
'%TESTUSER%' " -resolveSIDs:ON
```

Log Parser can be downloaded from this URL:

```
www.microsoft.com/downloads/details.aspx?FamilyID=
890cd06b-abf8-4c25-91b2-f8d975cf8c07&
displaylang=en
```

An unofficial Log Parser support site is available at www.logparser.com.

More Log Parser scripts

As you may have gathered, creating Log Parser scripts is an art in itself and there are books and courses on the subject. Because of that, we've provided some additional very useful SQL scripts to perform the functions indicated. These were specifically written by Dave Kleinman (a network-forensics specialist who has helped with this book) to target evidence of malware installation or intrusion. Because they're difficult to type in, they're also posted (ready to copy and drop) at the CastleCops Rootkit Revelations Forum:

These scripts are all single-line commands run from the command console. As usual, to open a command prompt, choose Start ⇨ Run, type **cmd** into the Open: field, and press Enter.

Any individual script listed here can be entered into a Notepad text file and saved with a `.sql` file extension. Then all you have to do to run the script at the command line is type in its filename — which saves you from having to input the entire script code just to execute the command.

For example, for the first command in the upcoming list of examples, you can insert the script code (located between the quotes) into a Notepad file, and then save that file as **process.sql**. Then, that script can be easily run by issuing the following command:

```
Logparser process.sql
```



Security auditing must be enabled and configured according to the specs described in Chapter 6, for these scripts to work.

- ✔ This command lists all processes that have ever started on the system — in descending order from the Security Event Log. (**Note:** It won't list processes stored in autoarchived logs, only those in the active Event Log.)

```
logparser "SELECT TimeGenerated,TO_LOWERCASE(EXTRACT_TOKEN(Strings,1,'|'))
AS processFileName,TO_LOWERCASE(EXTRACT_TOKEN(Strings,0,'|'))
AS ProcessID,TO_LOWERCASE(EXTRACT_TOKEN(Strings,2,'|')) AS
CreatorProcessID,TO_LOWERCASE(EXTRACT_TOKEN(Strings,3,'|')) AS
UserName INTO DATAGRID FROM Security WHERE (EventID IN (592))
GROUP BY TimeGenerated, ProcessFileName,Username, ProcessID,
CreatorProcessID ORDER BY TimeGenerated DESC"
```

- ✔ This command shows the total number of failed logon attempts and their respective IP, from the Security Event Log:

```
logparser "SELECT COUNT(EventID) AS
TotalLogonFailures,TO_LOWERCASE(EXTRACT_TOKEN(Strings,11,'|'))
AS SourceAddress INTO DATAGRID FROM Security WHERE (EventID
IN (529; 530; 531; 532; 533; 534; 535; 539;680)) AND
(SourceAddress IS NOT NULL) GROUP BY SourceAddress ORDER BY
TotalLogonFailures ASC"
```

- ✔ Here we output the same information as in the previous example, but we display it in bar-chart form instead:

```
logparser -o:CHART -chartType:Column3D -legend:ON -categories:ON
-values:ON -view:ON -chartTitle:"Host Logon Failure"
"SELECT TO_LOWERCASE(EXTRACT_TOKEN(Strings,11,'|')) AS
SourceAddress,COUNT(EventID) AS TotalLogonFailures INTO
FailLogon.gif FROM Security WHERE (EventID IN (529; 530;
531; 532; 533; 534; 535;539)) AND (SourceAddress IS NOT NULL)
GROUP BY SourceAddress ORDER BY TotalLogonFailures ASC"
```

- ✔ This command checks the last 50 pieces of software installed:

```
logparser "SELECT TOP 50 Value AS Product, LastWriteTime AS [Date
Installed]FROM HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\
Installer\UserData WHERE ValueName='DisplayName'ORDER BY
LastWriteTime DESC"
```

- ✔ This command shows the last time each service started:

```
logparser "SELECT Path, LastWriteTime INTO DATAGRID FROM HKLM\SYSTEM\
CurrentControlSet\Services\ WHERE ValueName='Start' "
```

- ✔ This command shows the last 50 files created on the C: drive:

```
logparser -i:FS "SELECT TOP 50 Path, CreationTime INTO DATAGRID FROM
C:\*. * ORDER BY CreationTime DESC"
```

- ✔ This command shows the last 50 files modified on the C: drive:

```
logparser -i:FS "SELECT TOP 50 Path, LastWriteTime INTO DATAGRID FROM
C:\*. * ORDER BY LastWriteTime DESC"
```

- ✔ This command finds application(s) that triggered Dr. Watson (a Windows debugger program that gathers error reporting information when an error occurs):

```
logparser "SELECT timegenerated, extract_token(strings, 0, '|') AS
WatsonError Into DATAGRID FROM system where sourcename=
'DrWatson' "
```

Checking Your System Resources

Another indication of successful malware intrusion is excessive consumption of memory, CPU, disk, or bandwidth resources. A rootkit can hide physical evidence of its presence on your hard drive (such as files, folders, and Registry entries) but the excessive drain it can place on system resources is hard to camouflage. Backdoor threats bring with them high bandwidth usage — particularly if your server has been set up as a database or server to host illegitimate activities.

Normally, rootkits try to be conservative in their resource requirements so as not to make their presence obvious. Nonetheless, there is often some noticeable effect on system performance — and resource consumption — if a backdoor is installed.

Malware processes are notorious for slowing down your computer and trying to grab a big slice of the CPU pie. Examining process activity with a number of different utilities (some rootkit specific, some not) can usually help you pinpoint the guilty party. We'll show you five different programs that you can use to do just that, in the “Examining active processes” section. Another resource that rootkit threats can put a drain on is hard-drive space. Again, disk usage varies with the threat installed and how much it requires for its

activities. If your computer or server is being used to host distributed illegal content, disk-usage requirements tend to go way up. Many backdoor trojans will download additional backdoors and totally compromise a host (imagine installing multiple doors in a wall) — and the resulting excessive resource consumption will be pretty obvious.

Checking your system resources can be a simple or complicated task, depending on the system and how familiar you are with it. Relatively inexperienced computer users may not know all the items and variables they should be looking at, even if they've been through this chapter. But the basics will steer you straight: Checking for activities and events produced by malware — including rootkits — need not be some gargantuan, complicated undertaking. With just a few simple tools (some already included with your operating system) and some freeware you can download from the Internet, you can check your system and maybe identify the causes of your problems.



For most of these checks, you have to use an Administrator account.

Matching activity and bandwidth

If your computer is suddenly very active while you're not doing anything with it, this is a sure sign that something or someone else is using it. You can monitor the activities on your computer to some extent using the Windows Task Manager Performance, Networking and User tabs. One easy way to open Task Manager is to right-click the desktop toolbar and select Task Manager from the pop-up menu or just use the good ol' (Ctrl+Alt+Delete) method. By checking these tabs, you can see what activity is occurring on your computer — and whether it's extreme (which is to say, possibly suspicious). Checking the User tab can show whether you have more users than before, and whether they're online.

To measure the bandwidth you're using in real time, first check the bottom of the Network tab in Task Manager. There you can see your Network Utilization given as a percentage. To measure your actual real-time bandwidth for sending (Tx) and receiving (Rx), you need a tool such as RxTx (which is freeware), available from the netfor2 Web site:

```
www.netfor2.com/RxTxPreview.html
```

To directly download the RxTx ZIP file, go here:

```
www.netfor2.com/RxTx32.zip
```

You need to check the activity and bandwidth usage to determine if your computer is doing a lot of communication, uploads and downloads you haven't authorized. If you know that normally your Internet-usage levels rarely exceed 5 percent, and find they've suddenly jumped to 50 percent or higher — and are staying there — you know something is wrong. If these levels are very high, *disconnect from the Internet immediately*. If your programs won't allow it, then pull the Internet connection plug out of the back of your computer. You will get an error message, but this effectively ends your Internet session — and disrupts whatever malware nastiness may be going on. Stay offline until you've determined the cause of the activity and corrected it.

Examining active processes

By scanning the process list you can quickly spot if any processes are using an inordinate amount of system resources, and then try to identify what the processes are. Malware processes are often big resource consumers. When you become more familiar with recognizing process names, you can distinguish if any process names are unfamiliar. If so, try Googling those processes or researching them, and make sure you are aware of a process's location on your disk — not just its name — before you start investigating.

In the following sections, we describe five programs you can use to view active processes: the Windows Task Manager, Process Explorer, IceSword, DarkSpy, and GMER (these last three are rootkit-detection programs). With the first three we've provided an image of the same process view using each of these programs. Because Process Explorer and IceSword use color coding to distinguish items in the display, this won't be evident in the images printed here but we've described how they're used.

Task Manager

Task Manager can be used to examine the process activity on your system. You can access Process Manager by hitting the (Ctrl+Alt+Delete) keys simultaneously on your computer keyboard. If you click the Process tab, Task Manager will display the active processes from all users.

The Process display features columns labeled as follows:

Column name	What it is
Image Name	This is the executable filename of the process that's running
User Name	This is the account under which the listed process was started

Column name	What it is
CPU	Shows CPU cycles consumed by the process
Mem Usage	Shows memory usage in KB for each active process

You can stop processes that have stopped responding by hitting the “End Process” button. The name of the process corresponds to the filename in the image column. Many malware files, particularly rootkit files and trojans, will try to mimic the names of legitimate operating system files. It’s very important to establish the location of any suspicious file before you make any final decisions about its threat potential. If you’ve confirmed a process as malicious, you need to end the process before you can delete the file from your hard drive. You will also have to be very certain you delete it from the proper folder location on disk. Unfortunately, the location of the file on disk (the complete file path) isn’t listed in Task Manager (see Figure 8-18), so you’ll have to search your disk to find the invader, or better yet, use Process Explorer, which identifies the file path for every active process. In Chapter 9, we go over methods of researching and identifying malicious files (and we reviewed some of those earlier in this chapter — such as searching online process/service databases and uploading the files to online antivirus scanning services).




Figure 8-18:
The active
process list
in Task
Manager.

You can use Task Manager to check system processes immediately, if a program hangs or odd behavior starts to crop up. Usually this utility can quickly help you identify the process causing trouble. It also gives you the ability to determine the filename of a process and to kill a troublesome process (by using the End Process button).



Keep in mind that nearly all rootkits hook the APIs that create Task Manager's active process display. This represents a real limitation when it comes to using Task Manager to examine your system for rootkit processes.



Task Manager won't show you any services loaded by active `svchost.exe` processes, or DLLs launched by `Rundll32.exe` — nor will it distinguish running processes from running services. Because most rootkits install kernel drivers (which Windows treats as services), being able to distinguish between running processes and services is helpful to rootkit analysis.

Process Explorer

Sometimes you have to take out the big guns — namely Process Explorer — to research process activity more deeply or see what services each instance of `svchost.exe` has loaded, or determine which DLLs `Rundll32.exe` has launched. Because Process Explorer can be run in combination with AntiHookExec, it can reveal both hidden process and service activity that Task Manager can't see. It can also check all the drivers installed on your system against a signature database to determine their authenticity — rootkit drivers, as of yet, haven't found a way to bypass digital-signature verification.



Process Explorer is included on your DART CD.

Chapter 9 provides images that show how to use Process Explorer to spot and stop the spurious Hacker Defender service, and how Process Explorer can check all resident drivers for digitally verifiable signatures. Process Explorer uses difference color highlighting to distinguish running services from running processes. Processes appear in lilac and services are in pink. It also allows you to stop running processes and services.

If a running process has an injected DLL or thread, you can see this by clicking the thread function in the Process Properties dialog box. Individual threads within a process can be killed by suspending the process and then killing the threads (then when the DLL is no longer active, its file can be deleted from your hard drive). All these functions are available from the Process Properties dialog box, which you access by right-clicking a process in the Process Tree and then selecting Properties. Figure 8-19 shows Process Explorer's active process list.

Process	PID	CPU	Description	Company Name
System Idle Process	0			
Interrupts	n/a		Hardware Interrupts	
DPCs	n/a		Deferred Procedure Calls	
System	4			
smss.exe	476		Windows NT Session Manager	Microsoft Corporation
csrss.exe	560		Client Server Runtime Process	Microsoft Corporation
winlogon.exe	584		Windows NT Logon Application	Microsoft Corporation
services.exe	636	1	Services and Controller app	Microsoft Corporation
svchost.exe	816		Generic Host Process for Win32 Services	Microsoft Corporation
sun1ThreatEngine.exe	1888		CounterSpy Threat Audit Engine	Sunbelt Software
SunProtectionServer.exe	1460		CounterSpy Active Protection	Sunbelt Software
svchost.exe	872		Generic Host Process for Win32 Services	Microsoft Corporation
spoolsv.exe	1204		Spooler SubSystem App	Microsoft Corporation
ewidoclit.exe	1688		ewido control	ewido networks
ewidoguard.exe	1716		guard	ewido networks
nod32krn.exe	1792	1	NOD32 Kernel Service	Eset
prevstsv.exe	1852		Panda Preventum+ service	Panda Software
loadsvcs.exe	360		PreEmpt Service Loader	PrivX Solutions, Inc.
vsmon.exe	296		TrueVector Service	Zone Labs, LLC

Name	Description /	Company Name	Version
SORTKEY.NLS			
sorttbls.nls			
CTYPE.NLS			
advapi32.dll	Advanced Windows 3...	Microsoft Corporation	5.01.2600.2180
comctl32.dll	Common Controls Lib...	Microsoft Corporation	5.82.2900.2180
dsname2.dll	dsname2	DiamondCS	2.00.0000.0900
dnsapi.dll	DNS Client API DLL	Microsoft Corporation	5.01.2600.2180
gdip32.dll	GDI Client DLL	Microsoft Corporation	5.01.2600.2818
hnetcfg.dll	Home Networking Co	Microsoft Corporation	5.01.2600.2180

CPU Usage: 100% Commit Charge: 41.70% Processes: 39

Figure 8-19:
The active
process list
in Process
Explorer.



Killing a legitimate process thread can wreak temporary havoc on your system (curable with a reboot), so be sure to verify that a thread or DLL is malware-related before doing that. Better yet, seek the advice of a security expert (via one of the Security Forums listed in Chapter 13) before attempting to use this approach on your own.

IceSword

IceSword can also show the Active Process list by clicking the Process icon under on the Function toolbar. IceSword will list hidden processes in red text, so it's an ideal way to target rootkit processes. IceSword uses red to distinguish hidden processes and services from visible components (which are listed in black). IceSword lists hidden services in red when you click the Win32Services icon located on the Function toolbar. We get deeper into using IceSword to detect hidden processes and services in Chapter 9, when we illustrate how IceSword can detect the Hacker Defender process and service.

Figure 8-20 shows you the IceSword active process list.

DarkSpy and Gmer

Like IceSword, DarkSpy has Process and Driver functions that list rootkit hidden processes and drivers in red. GMER has Process, Module (Driver), and Services functions. These functions list rootkit-hidden processes, modules (drivers), and services in red.




Figure 8-20:
The active
process list
in IceSword.



Both DarkSpy and GMER are both highly regarded anti-rootkit programs included on the *Rootkits For Dummies* DART CD. Both are discussed at length in Chapter 9.

Monitoring CPU cycles

Malware code that's badly or sloppily written will make its presence known, by slowing down your computer. Usage may show in the Windows Task Manager Performance Tab at 100 percent, and remain there for some time. However, not every instance of maximized CPU usage is due to malware. It could be an honest program that's incompatible with your particular system. If you've installed something recently, try disabling or uninstalling it to see whether the problem stops. If it doesn't, then it's time to check your process activity and bandwidth as explained in the previous section.

A good habit to get into is to open Windows Task Manager every time you log on to your computer. Then (and from time to time afterward), check your CPU levels to see how they're performing. If you suddenly notice a very high level of CPU consumption, either repeatedly or continuously, you should research the process responsible for causing it. You can use Process Explorer to determine the location of its source file on disk, and then submit it to one of the available databases we've recommended to determine its threat potential. You can also try booting into Safe mode to see whether the consumption levels come down.

To boot into Safe mode in Windows XP by using the System Configuration Utility, follow these steps:

- 1. Close all open programs.**
- 2. Click Start, choose Run, type `msconfig` into the Open: field, and then press Enter or click OK.**

The System Configuration Utility opens.

- 3. On the BOOT.INI tab, Check the `/SAFEBOOT` option and then click OK.**

Your computer will prompt you to reboot.

- 4. Restart your computer at the prompt.**

The computer starts in Safe mode.

- 5. When you're finished with troubleshooting in Safe mode, open `msconfig` again, and then (on the BOOT.INI tab) uncheck `/SAFEBOOT`, and click OK to restart your computer.**

When your computer is in Safe mode, run all your security scanners and use the diagnostic tools we've described throughout this book. You will be checking primarily for malware such as trojans, backdoors, key loggers, viruses, spyware, and adware. If a rookit is inactive in Safe mode, these conventional scanners can detect it and its payload. Dedicated rootkit scanners should be run in normal mode because they're programmed to look for hidden items; there is no advantage to running them in Safe mode, in fact, their authors don't recommend it.

If you can't correct the problems, contact a reputable security forum (see Chapter 13 for suggestions) — using a different computer if possible, so as not to spread any potential infections — and ask for help. It might be worth checking them anyway to ensure that you really have solved your problems. If security Web sites are not your thing, then call a local computer-repair shop or the service facility listed on your warranty card.

Chapter 9

Dealing with a Lying, Cheating Operating System

In This Chapter

- ▶ Realizing that it's not the OS's fault — it's been brainwashed
 - ▶ Checking out WinPE and Linux-bootable CDs
 - ▶ Reviewing some rootkit-detection tools
-

After you have all the basic security elements in place (as described in earlier chapters), you have some clear steps to take if you believe your computer is infected. This chapter details how to perform those essential tasks. If your computer is infected with a rootkit, you have to use tools with capabilities above and beyond those normally used to maintain good computer security. Often the only clue that a system is rootkit-compromised is negative: the computer shows infection symptoms but can't identify the cause — all traditional scanners and system-analysis programs are coming up empty-handed. That's when you have to dig deeper by using an assortment of dedicated anti-rootkit detectors and system-analysis tools. This chapter shows you how to do that.

There are many rootkit programs available — but unlike antivirus programs or anti-spyware programs, rootkit detectors vary in their ability to detect specific hiding techniques. That's why we suggest relying on not just one detector, but several different detectors when you're troubleshooting your computer; you get the best that each program has to offer. Several rootkit detectors (such as IceSword, GMER, and DarkSpy) offer suites of tools for detecting rootkit components; other anti-rootkit solutions excel at detecting just one item. Some detect hidden items but offer no removal capability; others include both removal and detection capabilities. There are also a few rootkit detectors that offer excellent detection but require a more advanced knowledge to decipher the results. While F-Secure (with its BlackLight scanner) and Sysinternals (with Rootkit Revealer) were pioneers in Windows rootkit detection from the very beginning, 2006 saw many antivirus companies release

rootkit scanners of their own — and that will probably be the tide of the future. Rootkit scanning will go mainstream by becoming part of general security solutions covered by antivirus companies. Still, it's likely that the separate dedicated scanners will continue to maintain their niche because of the superior solutions they offer.



Some more popular — and most effective — rootkit-detection programs include Rootkit Revealer (Microsoft Sysinternals), UnHackMe by (Greatis), DarkSpy (Mingyan Sun and Jianlei Shao), and GMER (Przemyslaw Gmerek). All these programs are included on your DART CD and are discussed in detail in this chapter.

Rooting Out Rootkits

As we mention in Chapter 1 (and detail in Chapter 7), rootkits hide their processes, files, and folders by using sophisticated hooking and filtering techniques. As a result, traditional methods of viewing the system state typically return no indication of foul play; the rootkit makes sure of that. As mentioned elsewhere in the book, system utilities such as Task Manager and Regedit (the Windows Registry editor) won't be able to expose the processes and Registry data that should (but doesn't) betray the presence of the rootkit. The lurking rootkit files won't be viewable in Windows Explorer — or even via the command line.

When a rootkit exists in this hidden state, we say it's *cloaked*. When the rootkit becomes *uncloaked* (that is, unhidden), the malware components it was hiding become uncloaked as well. To help you do that particular magic trick, you can use a rootkit scanner that's capable of *specifically detecting and identifying rootkits*. Different scanners use different strategies to accomplish this task:

- ✓ Some rootkit scanners can uncloak the rootkit so the malware's cyber-footprints become visible to traditional scanners.
- ✓ Some rootkit scanners can detect possible rootkit abnormalities but leave the interpretation of scan results — and rootkit removal — up to the user.



Rootkit scanners only provide the *means* of removing a rootkit; they don't remove it automatically. Instead, they wait for you to tell them what to do after showing you their results. That's because removing a kernel-mode driver can crash your system or remove some system functionality — and software developers don't want that liability. The programs that come closest to automatic rootkit removal are UnHackMe by Greatis and AVG Anti-rootkit

scanner by Grisoft — but they still request user confirmation before proceeding with removal.



It's very important to emphasize that once a rootkit is unclocked by rootkit utilities, the malware payload that accompanies the rootkit must then be removed — otherwise the payload can still proceed with its nastiness and may even reinfect the computer with the rootkit. Thus rootkit removal is a two-step process:

- 1. Uncloaking and removing the rootkit.**

Usually this step involves using special software tools that can find the rootkit and delicately “root” it out (sorry about that).

- 2. Removing the malware payload associated with the rootkit.**

This step normally uses conventional security programs such as antivirus, anti-trojan, and anti-spyware scanners. It may also involve manual deletion of stubborn rootkit components.

Cleaning a network

Ridding standalone machines of rootkits is challenge enough; if you're on a network, you have a more complex task. That's because when the host computer is compromised, data flow can allow the infection to spread across the entire network. Cleaning the individual client computers connected to a network server provides only a temporary solution. As long as the network server remains compromised, it can persistently reinfect the individual workstations. In addition, any client computer equipped with its own modem poses a potential risk to the entire computer network — especially if it has no active firewall. When infected, a client computer can pass the contamination back to the host — which can, in turn, compromise the remaining clients. Because data flow on a network runs in both directions (to and from individual machines), whether or not a rootkit is *persistent* (able to survive a reboot) becomes inconsequential; the server can always reinfect the client systems, and vice versa. The solution here is to unplug all connections to the network hub. This will sever the connections between all the networked clients, too, not just their connection to the server. Now each individual computer is isolated for the cleanup stage, which comes next. Then you must individually disinfect and reboot *all* networked computers — including the server — before any network connections are re-established. Provided all computers are cleaned appropriately, this systematic disconnection-and-disinfection strategy should effectively break the perpetual cycle of reinfection.

Before doing anything . . .

Because most rootkits alter kernel-level data structures (or even the kernel code itself), exercise caution before — and during — any attempt at rootkit removal. If you're not sure of what scan results mean — or you think you have a rootkit but are unsure of how to remove it — don't hesitate to ask for help at one of the online security forums (discussed in Chapter 13). The kernel is in a delicate balance. The authors of rootkit detectors are keenly aware of maintaining this balance when they develop their tools, but rootkit writers really don't care what happens to your system when you remove their handiwork. The wrong move could upset the kernel's balance; ask for help if you're unsure of how to proceed.

The best overall strategy

The best overall strategy for rootkit detection and removal is *not to rely on any single strategy*. The available tools use a variety of methods to detect rootkits, so be sure to use several of them to arrive at a correct diagnosis. Moreover, because scanners target rootkit behavior (rather than using signatures) scan results can be riddled with false positives or difficult to analyze. Plus, most detectors have specific strengths and weaknesses, and target different rootkit modifications. A diversity of solutions is available, so your best bet is to do a *cross-diff comparison* (to borrow a term from Rootkit Revealer) of the various scan results presented by different rootkit-detection tools: Combine several different detection tools to obtain a second, third, or even fourth opinion — and then coordinate and evaluate the results.

One thing is certain: When it comes to rootkits, nothing is a hard-and-fast science. There's no surefire method of detecting every rootkit out there, and examples of the species abound. Moreover, rootkit removal is a very delicate process that encompasses both removal of the rootkit *and* its accompanying threat.



Rootkit removal is a delicate, tenuous process, so it's best to proceed with the utmost caution and solicit qualified help if necessary.

Scanning Your OS from an External Medium

As we have discussed many times in the book, the integrity of the operating system is in question on a rootkit-infected host, which means an operating

system on a rootkit-compromised computer can't be trusted to return reliable results. As discussed in detail in Chapter 7, a rootkit intercepts operating system calls and filters the information the operating system returns to deny the rootkit's presence.

To combat the fact that rootkits cloak themselves from the host operating system, security experts developed a strategy to circumvent an operating system that has been brainwashed. Because one cannot trust the results returned by the operating system on a "rooted" (rootkit-infected) computer or network, it's helpful to leave the host operating system out of the equation — by viewing the target system from outside. You do this by booting to (and conducting scans from) an external storage device such as a CD-ROM, flash drive, hard drive, or the hard drive of another clean computer (slaving). When the scans of the infected computer are performed from an uncompromised operating system, the results *should* be trustworthy.

Unless you have a bootable external hard drive readily at your disposal, the easiest and most cost-effective method of capturing an external view is booting to an external CD-ROM. Booting to an external CD-ROM doesn't rely on the lying host operating system for its results; it relies on the untainted operating system on the CD to present an unbiased view of the infected target. Antivirus and anti-spyware scanners, when run from a bootable CD, can detect any malware present because unwelcome visitors won't be invisible to an external operating system the way they can be to the brainwashed OS of a rootkit-infected computer. Although even this method isn't totally foolproof, it's probably the most effective way to detect rootkits — particularly when resident rootkit detectors run on the suspect computer are turning up empty. Before declaring your computer "clean" when your instincts tell you better, scan the infected computer from an external operating system to get the final word on whether your system is infected.

Many bootable CD-ROM solutions exist, and sorting through them can take time. The following sections get you up to speed on what choices you have.

Microsoft WinPE

Microsoft WinPE — the Windows Preinstallation Environment — provides a basic set of tools for troubleshooting a Windows XP or Windows 2003 operating system from a bootable CD-ROM or USB drive. It loads only very basic services and drivers and a command-line utility; there is no graphical interface. WinPE isn't available to everyone; only Original Equipment Manufacturers (OEMs) and corporate clients have access to WinPE through the Windows Software Assurance program.

A few CD-ROMs can be purchased with the Windows Preinstallation Environment already installed, but these are rather pricey. That's because purchasing one essentially entails buying a licensed copy of the Windows operating system, in addition to the diagnostic tools that are provided.

Non-Microsoft bootable CDs

Fortunately, there are several free bootable CD-ROM build programs available online that present an alternative to using WinPE — and even offer expanded capabilities.



A successful WinPE alternative must have the capability to read and write to NTFS drives on Windows systems. Because a CD is only bootable if it contains a copy of an operating system, the most versatile of these CDs must be built using files from your own licensed copy of Windows, using either your Windows installation CD or your hard drive.

Windows-based CDs

We provide links to two types of free Windows CD build utilities and a bootable Linux CD. The Linux CD image can be downloaded without having to be built. That's because Linux is an open-source (free) operating system, so there are no licensing issues involved in obtaining a ready-made Linux CD image online to create a bootable Linux CD. However, bootable Linux CDs can only scan a folder at a time on NTFS partitions. If you compare Linux to Windows-based alternatives, there's a definite trade-off between the simplicity of creating a bootable CD and scanning versatility.

BartPE

BartPE or Bart Preinstallation Environment is probably the most popular build-it-yourself bootable CD for Windows. Bart PE is freeware (though donations are encouraged). BartPE's developer sponsors a Web site that provides complete instructions for building the BartPE CD:

www.nu2.nu/pebuilder/

BartPE provides read and write NTFS support — which means not only can malware be detected, it can be deleted from your hard drive. Bart PE has a graphical user interface that makes it easy to operate. To run third-party application programs from a BartPE CD, you must install plug-ins. Fortunately, you can get an extensive list of BartPE Plug-ins for security applications. The complete list of available BartPE Plug-ins (including a searchable plug-in database) can be found at the Bart Plug-in Repository at the following URL:

www.bootcd.us/BartPE_Plugins_Repository.php

BartPE offers many plug-ins that enable you to run security applications from your CD. For example, plug-ins exist for Ad-Aware SE, SpyBot-Search & Destroy, Avast Antivirus, Antivirus Personal Edition Classic, and McAfee Stinger. For alternative browser users, there is a Firefox plug-in available. And that's just for openers.

Before you build a BartPE CD, make sure your system meets these requirements:

- ✓ A Windows XP or Windows 2003 Installation CD or Windows installation source files on your hard drive
- ✓ 500MB of free disk space
- ✓ A read/write CD drive
- ✓ A copy of the most recently released version of PE Builder, which can be downloaded at the PEBuilder Web site: www.nu2.nu.



For BartPE newbies, a tutorial that gives precise information on installing plug-ins can be found here:

www.bootcd.us/PEBuilder_tutorial.php

Ultimate Boot CD for Windows (UBCD4Win)

The Ultimate Boot CD for Windows (or UBCD4Win) — another highly recommended freeware online CD builder — is an attractive alternative to the BartPE CD, simply because it does a lot of the work for you. This program has the BartPE Builder as its core, but it extends the basic BartPE offerings by including many well-known programs as part of the UBCD4Win's standard build.

There are many security programs, recovery programs, and backup programs included in the basic UBCD4Win CD build. This saves you the trouble and hassle of searching for and installing your own plug-ins, although the option for adding customized plug-ins is still available.

You can find a full list of the programs included with the UBCD4Win CD build at the developer Web site:

www.ubcd4win.com/

The Web site provides extensive help resources on building the CD and even a video tutorial to guide you along the way.

There's an online forum where you can keep up to date on new release information and post for help if you wouldn't mind a little assistance with creating the UBCD4Win CD:

www.ubcd4win.com/forum/

A rootkit-detection program called Rootkitty (included in the basic UBCD4Win build) automates a manual rootkit-detection method recommended by Microsoft. (We talk more about Rootkitty in the “Rootkit-Detection Tools” section later in this chapter.)

Linux-based CDs

The LinuxDefender Live! CD is a Linux-based bootable CD alternative. Instructions for creating this CD can be found at the BitDefender Web site:

www.bitdefender.com/bd/site/presscenter.php?menu_id=25&n_id=58

The LinuxDefender Live! CD comes with a copy of the BitDefender antivirus installed. Run from an external medium, it can scan both Windows and Linux-based systems for viruses. It has a graphical user interface and can disinfect (as well as detect) virus files. Because it operates from a Knoppix distribution of the Linux operating system, it can't scan your entire hard drive at once. All you can do is scan folders — but the advantage of using this tool is that the image is ready to burn to CD, and no CD build is required (as it is for Windows-based bootable CDs).



It's important to create a bootable CD-ROM before you actually need it — before your computer becomes infected — because it's not wise to ask the host OS of a (possibly) rootkit infected computer to participate in creating a troubleshooting CD-ROM. This especially applies to creating the Windows-based CD-ROMs, because Windows OS components from your computer become part of the fabric of the CD-ROM; they are, in fact, what makes the CD bootable. Transferring bits of a rootkit-infected OS to the bootable CD-ROM is not the way to go, for obvious reasons, so make and test one or more of the bootable CDs so you know you can depend on them when and if the time comes (though we hope it doesn't). If it's too late for that and your computer is already infected, make the CD on a clean system if possible.

File-System Comparison from Full Boot to Safe Mode

Malware (especially the rootkit variety) is designed to hide from virus and spyware scans when the computer is in Normal user mode, but sometimes you can catch it in the act if you're in *Safe mode* — a special diagnostic mode of Windows that's provided for troubleshooting purposes. (Usually, as you might expect, Windows runs in Normal mode.) A minimal set of system drivers (and very few Windows services and programs) run in Safe mode. Consequently, a rootkit and its kernel-mode driver (if it uses one) may not be running when you're in Safe mode.

A rootkit must be running (active) in order to hide itself and its payload; if it's not running, both are uncloaked — visible and susceptible to removal by traditional scanners that can run in Safe mode. Most antivirus and anti-spyware programs can be run in Safe mode. AVG Anti-Spyware (formerly ewido anti-malware), which is very effective at targeting trojans, can be run in Safe mode, as well. The system-analysis tools Autoruns and Process Explorer can also be run in Safe mode to spot malware and rootkit autostarts, drivers, and processes (providing the rootkit is not active). Autoruns and Process Explorer are included on your DART CD.

The point is that you may be able to zero in on and even solve rootkit problems without rootkit-detection tools at all; try running your anti-malware scanners in Safe mode and they may fix your problems without rootkit-detection tools. However, even though this result is theoretically accurate, we still strongly suggest confirming that your system is clean by running a couple of rootkit scans in Normal mode afterwards.

If your scans in Safe mode reveal nothing despite ongoing symptoms, your computer or network may be compromised by a rootkit that runs in Safe mode. (Yes, unfortunately, they exist — and that's actually the norm nowadays.) To determine whether this is the case, you should perform scans in normal Windows mode, with the rootkit programs included on your *Rootkits For Dummies* CD: DarkSpy, GMER, Rootkit Revealer, and UnHackMe — all of which are discussed in the “Rootkit-Detection Tools” section later in this chapter.



The Apropos rootkit was one of the most commonly encountered commercial-adware rootkits until its distributors voluntarily withdrew it in late 2005. Occasionally, we still see some users infected with Apropos (there's usually quite a delay between the time a threat ceases to be distributed and the time you stop seeing it on users' systems). Apropos is also an example of a rootkit that doesn't run in Safe mode; its Registry autostarts (which launch its randomly named processes) suddenly become visible in Safe mode. Apropos also installs a hidden spyware service that's visible in Safe mode. HackerDefender, on the other hand, is a rootkit that *isn't* visible in Safe mode. It inserts entries under the `HKLM\System\CurrentControlSet\Control\SafeBoot` Registry key to guarantee that its driver (services) will run in Safe mode.



We highly recommend you run a cleanup program (for example, CCleaner) before you run your scans, and then scan twice — first after you've booted up normally, and again after you've shut down your system and started up again in Safe mode.

Checkpointing Utilities with Offline Hash Databases

It's impossible to distinguish whether a system file of the proper size and name is truly authentic by merely inspecting a directory listing. Even checking file properties is not always enough; the file size and timestamp may have been doctored to match those of the authentic file. Even the version information can be doctored so the file appears to bear the Microsoft label. For all intents and purposes, the file can appear to be the real thing — even if it's anything but.

There is a relatively simple method of determining whether a user-mode rootkit has replaced any critical system files with its own trojanized versions that bear the same names and/or property information. This method can help assess the extent of the damage a user-mode rootkit may have inflicted on the operating system. It uses an indicator called a *checksum* or *hash value* — the unique digital footprint of a file — to determine whether system files have been replaced or tampered with.



Each file can be analyzed to produce a checksum, which is basically a digital signature calculated by applying the same predetermined mathematical formula (known as an *algorithm*) to every file in question. Some popular checksum algorithms are MD5 and SHA-1. Not only do the MD5 and SHA-1 algorithms reduce a file's unique footprint to a relatively small string of digits, they also encrypt the checksum data to make it more secure.

Using cryptographic checksums to assess file integrity is a comparative method that requires you to take a cryptographic system “before” snapshot (a baseline, as described in Chapter 10) of critical system files on a clean — preferably newly installed — operating system. The “before” snapshot represents the gold standard against which all future checksums or hashes will be compared. To ensure the integrity of these clean checksums, make sure they're stored on external media so the data can't be tampered with. Some rootkits are known to intercept the reporting of incorrect hash values to antivirus software as a way to thwart detection. Storing the data externally introduces an effective countermeasure to ensure safety.

Another rootkit trick is changing file permissions on critical system files. Analyzing whether security permissions have been altered is an operation incorporated into checksum comparisons.

Verifying files with FileAlyzer

The developers who gave us Spybot-Search&Destroy, a highly regarded anti-spyware program (included on your DART CD), also created FileAlyzer.

FileAlyzer reveals a lot of extra information about a file that far exceeds the information available in the Windows Properties dialog box. It's easy to get a quick FileAlyzer analysis of any chosen file because it inserts a clickable option into a context menu. Simply right-clicking a file in Windows Explorer and selecting Analyse File with FileAlyzer from the context menu allows you to see the identifying file information.

Okay, here's an example that actually happened to one of your authors in real life. Let's say you recently noticed an unknown file called `ntpmon.exe` on your desktop and decided to research its origin.

First, you right-click the file and chose Properties to bring up the Windows Properties dialog box for the `ntpmon.exe`. No dice; Figures 9-1 and 9-2 illustrate how you can't decipher anything about the program's source from the information provided by the Windows Properties dialog box.




Figure 9-1:
The
Windows
General
Properties
dialog box.

Clicking the Summary Tab also comes up completely empty, as is illustrated in Figure 9-2.

So then you decide to use FileAlyzer to determine whether it could expand upon the information provided by the Windows Properties dialog box. Figure 9-3 shows what you learn about `ntpmon.exe` by right-clicking the file and selecting Analyse file with FileAlyzer from the context menu.




Figure 9-2:
The
Windows
Properties
summary.

By clicking the hex-dump option, you can examine the strings embedded in the file. These strings, or readable text characters, are listed under the column labeled Text. The strings displayed enable you to locate the identifying information you need in order to determine the origin of the file. Fortunately, the strings clearly identify the heretofore-unknown executable file as Sysinternals Process Monitor, authored by Bryce Cogswell and Mark Russinovich. (Whew.)




Figure 9-3:
FileAlzyer
scan.

You can also see that files have created, modified, and accessed dates in their properties. All these dates are very meaningful for evaluating and assessing files that may have been corrupted by a hacker or malware infection (though unfortunately, some really devious malware exists that can change these dates; almost nothing is sacred when you're dealing with highly motivated and clever malware coders). Notice that you've obtained this information without opening or altering any dates associated with the properties of the file.

If this scenario sounds intriguing, you can download FileAlyzer for free from the Safer-Networking Web site:

```
www.safer-networking.org/files/filealyz.exe
```

Verifying file integrity with other utilities

Here's a small sampling of the more popular file-and-system-integrity monitors available online.

- ✓ **AccuHash 2.0:** (Freeware) AccuHash 2.0 is a Windows utility that uses checksum verification to ensure the accuracy of data files. It supports CRC32, MD5, and SHA-1 checksum algorithms. For more information and to download AccuHash 2.0, please refer to this link:

```
www.accuhash.com/news-20060209.html
```

- ✓ **Sigcheck (Freeware):** If you haven't yet made a backup copy of all critical system files for a cross-comparison, this tool may be of some help. Sigcheck is developed and distributed by Microsoft Sysinternals. Here's where to get it:

```
www.sysinternals.com/Utilities/Sigcheck.html
```

Sigcheck is a command-line program that you can use to verify that file images are digitally signed; you can also use it to retrieve file-version information. An especially nice feature: You can specify the directory that you want Sigcheck to analyze, and even specify search parameters. For example, if you want to create a list of all unsigned executable files in the Windows system folder (not a bad idea), you can specify that by using the following command (assuming C: is the drive where your OS is located):

```
sigcheck -u -e c:\windows\system32
```

- ✓ **Tripwire (Commercial):** Tripwire is perhaps the most widely known and popular integrity scanner out there, providing thorough detection and protection to standalone and network server computers. Notifications to administrators or IT personnel can be provided via e-mail or pager.

```
www.tripwire.com/products/index.cfm
```

- ✔ **Data Sentinel (Commercial):** Data Sentinel has a simple and easy interface, with controls for doing fast or normal checks on files. It scans over 55 MB of data per second, monitoring 15 file and seven Registry properties, and can be configured to e-mail the results.

www.ionx.co.uk/html/products/data_sentinel/index.php

- ✔ **MD5summer (Freeware):** The program creates MD5 cryptographic checksums of files on Windows computers. It's your basic file-integrity scanner (as described previously), with no frills. You can create a baseline with it for later comparisons. MD5summer is a good standalone product for single computers.

www.md5summer.org

Rootkit-Detection Tools

The many rootkit-detection tools can be grouped according to the methodologies they use to detect rootkits. Some tools are based on a cross-difference comparison of your hard drive that employs different scanning techniques. Others use signature detection. Most detect rootkits by looking for their behavioral characteristics; they take a *heuristic* approach (see the nearby Technical Stuff icon for a definition). Others rely on inside-the-box-to-outside-the-box comparison scans of the host computer. Still other tools target the hooking of the data structures referenced by user processes or by the operating system. Some tools are comprehensive and can detect hidden processes, files, ports, drivers, and Registry entries of both user-mode and kernel-mode rootkits. Others concentrate solely on hook detection. Many rootkit utilities detect and list rootkit changes but don't provide any means of removing the rootkit. A few provide both detection and removal capabilities.



Heuristics is a technique that uses rules that have been formulated by identifying the common behavioral characteristics associated with a class of threat. These rules are then used to create an algorithm (solution) that can be applied to determine the *likelihood* that an entity constitutes a threat (as opposed to using a database of existing threat signatures like most conventional scanners do). If the rules are sound, then detection should be accurate. If the rules are too broad or non-specific, the net will be too large — and false positives may crop up in the scan results. Some heuristic scanners compare their scan results against a *white list* of benign programs to weed out false positives. Most anti-rootkit programs are heuristic, as opposed to signature-based; they rely on their ability to detect behavior typical of rootkits (such as modification of data structures or kernel objects in memory) to produce their results.

Symantec has compiled a list of the top 19 threats that exhibit rootkit characteristics. The list indicates that 17 of the 19 threats employ hooking; the remaining 2 rely solely on more sophisticated DKOM (Direct Kernel Object Manipulation) techniques. Seventy-five percent of the threats hook kernel data structures such as the System Service Descriptor Table (SSDT) or the Interrupt Descriptor Table (IDT) by installing a kernel driver. Because 90 percent of the top stealth threats employ some sort of hooking, using tools that base their methodology on hook detection should catch the majority of rootkits. Rootkit detection has evolved to the point where most rootkits that install a kernel-mode driver SYS file and employ system hooking are detectable by the currently available anti-rootkit tools. However, as discussed in Chapter 7, the rootkit landscape is ever-evolving; the new challenge for rootkit writers will be to make these current tools obsolete. The most sophisticated rootkits attempt to thwart current detection techniques, and hooking will inevitably be replaced by DKOM and inline kernel hooking as the preferred method of rootkit subversion.

Rootkit detectors also vary in their user-friendliness. They range from the newbie-friendly UnHackMe to those that produce only a cryptic scan report that even the most experienced user would have trouble interpreting.

One of the most comprehensive rootkit-detection-and-removal programs available today is IceSword. Not only does it detect processes hidden by DKOM modification of the process list, but it provides the facility to remove rootkits within the program. Because of that, we spend quite a bit of time discussing it in this chapter. Two relative newcomers to the anti-rootkit arsenal — DarkSpy and GMER — also take a multi-functional approach; in some cases, they've even surpassed IceSword's detection capabilities. Having been introduced more recently, they incorporate detection of nearly all known rootkit-hiding techniques that preceded their release. They have both been updated since their initial releases to expand their detection capabilities and improve their functionality.

BlackLight and UnHackMe are the easiest of these tools to run because they provide automatic detection-and-removal capability. They uncloak the rootkit and its rooted files so they're visible for removal by standard security programs (or manual methods). Although Rootkit Revealer is an excellent rootkit-detection tool, it provides no means of removing a rootkit. IceSword and DarkSpy are highly effective total solutions geared toward intermediate to advanced users; GMER is a very effective detector that offers some user-friendly functions that beginner or intermediate users should feel comfortable using. AntiHookExec allows system utilities such as Task Manager to return accurate and unbiased results, even if a user-mode rootkit is in residence.



When you're running rootkit-detection tools that perform an intensive full-system scan (such as Rootkit Revealer, BlackLight, and others discussed in the following sections), minimize false positives by leaving your system as idle as possible. All activity on your computer should be suspended: close all open Windows and programs, and disconnect from the Internet to prevent background programs from auto-updating during the scan. Disable active protection and any scans scheduled by security programs that may engage during the analysis. Refrain from using your computer while running the anti-rootkit scan; don't touch your mouse or keyboard or it will interfere with the results. Because most rootkit detectors launch services to scan your system, your protection programs may complain. When Nancy ran a BlackLight scan on her system (for example), WinPatrol's Scotty barked up an alert, which resulted in a false positive in the scan results. Make sure Scotty and other active protection programs are turned off to minimize any false positives that might be generated by active-protection notifications.



Currently available rootkit tools can detect all known public rootkits out there today. However, because the tools target different rootkit techniques, it's important to combine them for complete and effective coverage. Running BlackLight and Rootkit Revealer together targets both those rootkits that use API hooking to hide and those that use FU's method of hiding processes (using DKOM). Rootkit Revealer shows you Registry entries in addition to hidden files. BlackLight can detect processes hidden by DKOM modification of the Process List, but Rootkit Revealer cannot. IceSword is like having a suite of system-analysis and rootkit-detection tools all under one roof. IceSword can detect most known public rootkits today but you need to be able to interpret the results and know how to effectively disarm the rootkit. GMER and DarkSpy have improved upon IceSword's detection capabilities by targeting the most sophisticated inline hooking and DKOM hiding techniques used by the current crop of rootkits. These last two programs provide multiple functions and include rootkit-removal capability.



Scanning with an updated antivirus and anti-spyware program is an advisable follow up to using any rootkit-removal utility. Additionally, be sure to do a system assessment to research the extent of rootkit-caused collateral damage.

Autoruns: Aiding and abetting rootkit detection

Autoruns (provided on your DART CD) is useful in investigating startup program-launch points. Autoruns is not a rootkit-detection tool per se, but rather, a system-analysis tool you can use to detect persistent rootkits.

Autoruns lists nearly every location on your system that starts programs when your system boots. These locations are either Registry points or folder launch points on your hard drive. Autoruns has a useful file-comparison feature, which can target launch points for the rootkits that don't run in Safe mode. To take advantage of Autoruns file comparison feature, you must first use another Autoruns feature that enables you to save a log of your computer's autostart locations.

To test for rootkit autostarts, you should first run an Autoruns scan in normal Windows mode and save the log of all your system's autostart entries by choosing the File ⇨ Save menu option. Next, you need to boot into Safe mode and choose the Autoruns File ⇨ Compare feature. This compares the Safe-mode scan results to those you just saved in the scan log — and generates a difference log.

Assuming the rootkit does not run in Safe mode, the difference log lists any hidden rootkit autostarts that can be seen in Safe mode (when the rootkit was not running) versus those seen in Normal mode when it was. The image data for these files is listed under the Image column of the Autoruns main display. The image indicates the file path of the program (executable file) on disk that's being started by each flagged autostart entry.

Because all persistent rootkits have autostarts that guarantee the rootkit runs at system startup, using Autoruns is an excellent method to discover rootkit startups. Moreover, kernel-mode rootkits have additional autostarts for the services and drivers they load. Okay, we're not suggesting that Autoruns has to be the only tool used to detect rootkits, but it can detect hidden rootkit autostarts for rootkits that don't run in Safe mode — and do that in a matter of minutes. It can even detect user-mode rootkit autostarts in Normal mode. (We detail that capability a bit more in the "Using AntiHookExec with Autoruns" section later in this chapter.)

Rootkit Revealer

Another very valuable Sysinternals tool — developed by the authors of Autoruns, that's specifically dedicated to rootkit detection — is Rootkit Revealer. Rootkit Revealer (RKR) can detect files, folders, and Registry keys that are hidden due to the hooking of user or kernel-level (native) APIs.



Because Rootkit Revealer's detection method relies on hook recognition, it cannot detect rootkits that achieve their stealth solely through DKOM. Your best bet is to run this program in combination with another rootkit-detection program (such as DarkSpy, GMER, IceSword, or BlackLight) that does incorporate DKOM detection. One of the features that we like best about RKR is

that it lists rootkit autostart locations in its scan report. Knowing a rootkit's Registry launch point can greatly simplify rootkit removal. Rootkit Revealer provides no disinfection functions for removing rootkits, but it does enable you to save a log of the scan results. You can post your RKR log to one of the security sites listed in Chapter 13 and get assistance with interpreting it.

How Rootkit Revealer works

RKR's approach relies on the known fact that most rootkits hook various data structures (namely the IAT, EAT, SSDT, and IDT) to hide rootkit files. Because of this, the operating system of a rootkit-infected computer will return results generated by hooked APIs that can't be trusted (as discussed in detail in Chapter 7).

Because many rootkits have targeted the Rootkit Revealer executable file, RKR's developer has crafted a clever random-renaming technique to counteract the rootkits' evasion tactics: Each time a RKR scan is performed, RKR is launched from a new, randomly named copy of the RKR's executable file. Because the rootkit-detection program has a new name each time it runs, the rootkit can't identify the RKR process. So if you notice a new randomly named process running from your temp folder while you're performing a RKR scan, there's no cause for alarm.

Because Rootkit Revealer scans the hard drive for physical evidence of rootkit tampering, RKR's focus is on targeting persistent user-mode and kernel-mode rootkits. Rootkit Revealer uses a cross-view comparison method to target possible rootkit anomalies. It performs a higher-level file-system scan using normal Windows API calls to generate its results. It also scans the Registry and file system at a very low level that circumvents using the higher-level Windows APIs that cannot be trusted to return legitimate results. Rootkit Revealer then does a cross difference comparison between the results generated by the higher-level APIs and those generated by using low-level (raw) disk reads. The discrepancies in the results between these two types of scans are then flagged and presented in a log that can be saved for evaluation. The lower-level scan reads the NTFS Master File Table (MFT), raw Registry hives, and the NTFS and FAT file-system data structures. The higher-level scan obtains its results by using a DOS directory command that makes use of conventional Windows APIs.

A very useful feature of RKR is its ability to scan remote clients when launched via the Sysinternals command-line utility PsExec. The command to accomplish this is

```
psexec \\remote -c Rootkitrevealer.exe -a c:\windows\system32\Rootkit.log
```

This command stores the contents of the RKR log in a file called `Rootkit.log` located in the `system32` folder on the host machine.

The `\\` options means that Rootkit Revealer should be run on all computers in the current domain. Another alternative is to specify the target computers you would like to scan in a text file that's then passed as an argument to the `psexec` command. To see more about the usage of the PsExec utility (and how you can use it to execute applications on remote computers), please refer to the following link:

www.sysinternals.com/Utilities/PsExec.html

Interpreting Rootkit Revealer results

The scan results note discrepancies by using any of these four alerts:

Visible in Windows API, directory index, but not in MFT.

Visible in Windows API, but not in MFT or directory index.

Visible in Windows API, MFT, but not in directory index.

Visible in directory index, but not Windows API or MFT.

Rootkit Revealer is excellent tool for detecting both user- and kernel-mode rootkits that use hooking — but it doesn't offer any method of removing any files or Registry keys it flags in its scan results. The results must be interpreted and assessed for false positives. Any activity on your system that occurs while you're conducting a scan produces entries in RKR's results. Some of the more common causes of such activity include these:

- ✓ Background program activity such as
 - Auto-updating
 - Scheduled scanning
 - Active-protection alerts
- ✓ Keyboard or mouse movement
- ✓ Windows writing to System Volume information (for System restore)



Users of the since-replaced Kaspersky Antivirus version 5.0 who had the stream function enabled may have noticed that RKR listed many files that use alternate data streams (ADSs) in its results. This happened if the Hide NTFS Metadata Files option wasn't checked in RKR. Kaspersky Antivirus 5.0 used

the ADS feature of NTFS files to store file checksums so it could operate more efficiently. By quickly comparing a current file checksum to the one it had previously stored in the ADS portion of a file, it could determine whether a file's integrity had been compromised. If it hadn't been, then the program skipped a deep scan of that file and moved on. What seemed like a good idea for one program had an undesirable side effect on another: RKR listed all those legitimate ADS files in its scan results.

Files that use ADS streams are visible to Windows, but the data in the ADS portion of the file won't contribute to the file's overall size. Therefore the original file is visible in Windows Explorer, but the ADS data (which is listed after the colon) is invisible — and its size isn't added to the original file size. RKR shows this alert for ADS files: `Hidden from Windows API`. ADS files aren't necessarily harmful, and as outlined above, the ADS feature of NTFS files can be used for legitimate purposes. Kaspersky has since released Version 6.0 of its antivirus, and this newer version has abandoned using the ADS file feature to improve its scanning efficiency. Therefore Kaspersky antivirus users are no longer confronted with a mass of flagged-but-innocent files when they do a RKR scan. They can see malicious RKR ADS entries more easily, and wind up with a smaller RKR log.

Users of Alcohol and Daemon Tools will see entries generated from those programs in the RKR scan results.



The following Registry key is often encountered as a false positive; perhaps belonging to some hidden-but-benign Registry data. Though its origin remains mysterious, the entry is not malicious:

```
HKLM\S\C\webcal\URL Protocol
```

Knowing where to seek help

If you do find some consistent anomalies, then research the results. Most, but not all, rootkits will have been encountered already; consequently, many have removal directions posted online. If you're at all uncertain or reluctant to attempt removal yourself, use the list of security forums we provide in Chapter 13. You can also Google the results of any flagged files as a way to research removal methods. Be sure to double-check the directions too. Sometimes new variants emerge, or more effective removal directions develop that supersede the original removal instructions in their effectiveness.

Chapter 13 refers you to some highly regarded online security forums that can provide assistance with the interpretation of your rootkit-scan reports. These include the Rootkit Revealer Forum at Sysinternals, which you can find at the following URL:

```
www.sysinternals.com/Forum/default.asp
```



More information on Rootkit Revealer can be obtained in the Help function provided in the program and at the Sysinternals Web site:

www.sysinternals.com/Utilities/RootkitRevealer.html

F-Secure BlackLight Beta

BlackLight by F-Secure is a rootkit scanner that does an excellent job of ferreting out hidden rootkit processes, files, and folders but not hidden Registry keys. Many of the files detected by RKR will also be flagged by a BlackLight scan, but BlackLight enables you to rename the flagged files — which deactivates the rootkit and makes what it was hiding visible, after a reboot. Don't let BlackLight's Beta designation scare you; this program is tried and trusted — it's been around since Windows' rootkits first gained public prominence. BlackLight can detect rootkits that use DKOM techniques to modify the active process list (PsActiveProcessList) — however, it cannot detect the more advanced method of hiding processes introduced by FUTO that involves PspCid Table modification (the kernel data structure that keeps track of both processes and threads). Scanning with BlackLight is simple and fast — for that reason, it has become one of the most widely used tools to detect and disable rootkit processes.

You launch BlackLight by double-clicking its executable file (`blbeta.exe`) from its installed location. Figure 9-4 illustrates the BlackLight program display. You initiate a scan by clicking the Scan button, and BlackLight scans your system. When it's finished, it reports any hidden files it has found. In Figure 9-4 you can see how the BlackLight display looks after hidden files are detected:





Figure 9-4:
Scanning
for hidden
items with
BlackLight.

Clicking the Show All Processes button updates the display to show you a list of all hidden items found — and gives you the option of renaming any files it has listed. Figure 9-5 shows a list of hidden files found on a system infected with HackerDefender (HxDef variant). A hidden process called `control.exe` is listed but it's not malicious; it actually belongs to Sandboxie, a virtualization program included on your DART CD. The `bdc1i100.exe` file, on the other hand, is a process related to HackerDefender's backdoor port function.

Figure 9-5:
Cleaning
hidden
items with
BlackLight.



Note that BlackLight doesn't remove rootkits. If there are hidden files, you can rename them to prevent them from loading when you reboot; this causes the rootkit to become inactive, after which, you can remove the renamed rootkit files.



Use extreme care here — and rename *only* enough files to disable the rootkit. This will minimize any unintended side effects. For example, if a rootkit driver that's typically a .SYS file, or a rootkit executable (.EXE) file, turns up on the list of hidden files, then rename *only* those items — but don't rename all the files and folders identified by the scan (rootkits sometimes hide legitimate system files; renaming too many files can have a negative impact). After renaming a minimal amount of files, reboot your computer, and the other hidden malware components and any renamed items should be visible for removal. When you've killed off all the malware components you can find, do a repeat BlackLight scan to verify that hidden threats no longer exist.

After you reboot, any installed rootkit service will be visible in the list of services displayed in the Services Console (`services.msc`). The rootkit service should be disabled and its file deleted. Any other files associated with the rootkit must be deleted manually — after which, be sure to follow up by running conventional antivirus and anti-spyware scanners.

The BlackLight Beta creates a log file `fsbl-<date-and-time>.log` in the same directory as its program executable file (`blbeta.exe`).



At the time of this writing, F-Secure is still offering the BlackLight Beta as a free download. More detailed instructions, download information, and a comprehensive tutorial can be found on the F-Secure Web site:

www.f-secure.com/blacklight/help/

It's advisable to run both Rootkit Revealer and BlackLight and cross-correlate your results. RKR will also detect hidden Registry entries, and if you're unsure about the scan results, cross-comparing the scan reports may help you clarify them.



BlackLight can reverse DKOM modification of the Process List — which makes hidden processes visible. Note, however, that it can't detect processes that were hidden by modifying the PspCid table, nor can it detect drivers that were hidden by modifying the Loaded Drivers list in kernel memory. (Both of these super-sneaky hiding methods were introduced by FUTO.)

IceSword

IceSword can detect hidden files, folders, processes, and services. Unlike most other anti-rootkit programs, it can also detect hidden autostarts, browser helper objects (BHOs), Registry entries, and Windows message hooks. IceSword enables you to disable and eliminate the rootkit from within the program without requiring a system reboot for changes to take effect. It's considered an advanced tool, but its color coding makes identification of rootkit components relatively easy.



Because IceSword's SSDT function may display “innocent” programs in red be *very* careful about what you decide to delete.

Detecting rootkit changes with IceSword

One of the properties that makes IceSword so useful is that it provides the means of stopping, disabling, and removing the rootkit from within the program. Many people classify IceSword as a tool for intermediate to advanced users. However, because it presents information in a simplified format, the results are easy enough for even a quick beginner to interpret. IceSword uses a red-and-black color-coding scheme to differentiate the items it lists. Black items usually indicate trusted entries; red items, depending on which IceSword function is being used, can indicate either hidden or suspicious entries that merit further investigation.



A *process or service* listed in red indicates the item is hidden and is a definite indicator of rootkit activity. Therefore, using IceSword's Process and Win32 Services functions to reveal hidden processes and services (as shown in red) enables you to detect rootkit files. Clicking Kernel Modules displays kernel-level drivers. Unfortunately, no color coding is used to distinguish hidden drivers from the legitimate ones; IceSword displays *all* drivers in black text.

Figure 9-6 illustrates IceSword detecting the HackerDefender process `hxDef100.exe`, using the Process function.

ImageName	PID	PathName	Base Priority	EPROCESS	SI
System Idle ...	0	NT OS Kernel	0	0x00559580	Re
System	4	NT OS Kernel	0	0x812ED030	Re
IceSword.exe	208	C:\Tools\IceSword_en1.12\is_en\iceSword.exe	8	0x8119A5E8	Re
wuauclt.exe	264	C:\WINDOWS\system32\wuauclt.exe	8	0x81151030	Re
lview32.exe	268	C:\Program Files\lrfarView\lview32.exe	8	0x81132D40	Re
Control.exe	284	C:\Program Files\Sandboxie\Control.exe	8	0x81133C30	Re
smss.exe	368	C:\WINDOWS\system32\smss.exe	11	0x811DD898	Id
Start.exe	396	C:\Program Files\Sandboxie\Start.exe	8	0x81133668	Re
csrss.exe	428	C:\WINDOWS\system32\csrss.exe	13	0x81184030	Re
winlogon.exe	452	C:\WINDOWS\system32\winlogon.exe	13	0x81197640	Re
services.exe	496	C:\WINDOWS\system32\services.exe	9	0x81188698	Re
lsass.exe	508	C:\WINDOWS\system32\lsass.exe	9	0x81181328	Re
svchost.exe	676	C:\WINDOWS\system32\svchost.exe	8	0x81167440	Re
svchost.exe	760	C:\WINDOWS\system32\svchost.exe	8	0x81153030	Re
svchost.exe	812	C:\WINDOWS\system32\svchost.exe	8	0x8113F440	Re
svchost.exe	892	C:\WINDOWS\system32\svchost.exe	8	0x8116C030	Re
svchost.exe	948	C:\WINDOWS\system32\svchost.exe	8	0x811245E8	Re
explorer.exe	1152	C:\WINDOWS\explorer.exe	8	0xFFBA11C8	Re
spoolsv.exe	1164	C:\WINDOWS\system32\spoolsv.exe	8	0xFFBA11B8	Re
hxDef100.exe	1380	C:\Documents and Settings\swat\My Documents\hxDef100\hxDef100.exe	8	0xFFB53030	Id
VMwareService.exe	1472	C:\Program Files\VMware\VMware Tools\VMwareService.exe	13	0xFFB40708	Re
VMwareTray.exe	1508	C:\Program Files\VMware\VMware Tools\VMwareTray.exe	8	0xFFB59618	Re
VMwareUser.exe	1568	C:\Program Files\VMware\VMware Tools\VMwareUser.exe	8	0xFFB5CAE8	Re
alg.exe	2028	C:\WINDOWS\system32\alg.exe	8	0x8117E608	Re

Figure 9-6:
IceSword
Process
scan.

Figure 9-7 illustrates IceSword detecting the HackerDefender service `hxDef100.exe`, using the Win32 Services function.



When it comes to the SSDT function, a red entry doesn't *always* confirm a rootkit association. For example, if SSDT is selected to show hooked entries in the System Service Descriptor Table, there will most likely be a few red entries listed that can be attributed to your antivirus or other legitimate security programs. In fact, any entries that map to any module other than the operating system kernel (`ntoskrnl.exe`) will be displayed in red text.

When the SSDT function is selected, it's easy to verify the SSDT has been hooked by inspecting the `Current Addr` and `Original Addr` columns of red entries. Notice that these addresses will always be different for hooked (red) entries. In Figure 9-8, hooked entries are listed in light gray text, and the `Current` and `Original Addr` for these items differs (in this case, the items in red belong to a legitimate virtualization program — Sandboxie — which is included on your DART CD).

Figure 9-7:
IceSword
Win32
Services
scan.





Figure 9-8:
A System
Service
Descriptor
Table scan
in IceSword.



It's best to use IceSword's Process and Win32 Services functions first, to confirm the presence of the rootkit. Then you can use the Kernel Modules (drivers) and the SSDT functions to corroborate the information obtained using the Process and Win32 Services functions. If a rootkit is detected (and illegitimate, red-coded entries turn up after using the SSDT function), you can use the SSDT function to assess the effectiveness of the removal; just run that function again, after rootkit removal is complete.

Interpreting IceSword results

The following pairs represent some legitimate programs (and their associated drivers) that create red entries in the SSDT:

- ✓ **SpySweeper:** `ssi.sys`
- ✓ **Prevx:** `pxfsf.sys`
- ✓ **Process Guard:** `procguard.sys`
- ✓ **Panda Antivirus:** `shlddrv.sys`, `PavProc.sys`
- ✓ **Daemon Tools:** `-D347bus.sys`
- ✓ **Sandboxie:** `sandbox.sys`
- ✓ **Kaspersky Antivirus:** `klif.sys`

This is by no means an exhaustive list. Any information listed in the SSDT should be confirmed by viewing the entries in the active process list (which can be viewed by clicking Process) and the service list (which you can view by clicking Win32 Services). These functions display only hidden items in red. Hidden items are a stronger indication of rootkit activity than is SSDT hooking alone.



Be very careful to Google any flagged files to acquire more information before taking any removal actions. Another way of checking an item is by seeing if you can locate the responsible file or Registry entry using Windows functions, like Windows Explorer, Task Manager and the Windows Services Console. If Windows cannot see the file or Registry entry — and you have Windows set to display hidden files and folders — then it's a pretty safe bet that the suspect file represents a hidden rootkit component.

When a rootkit service or process is stopped, the rootkit will be uncloaked and visible to Windows, though you may have to enable Viewing of Hidden Files and Folders. It's possible that a rootkit file may have its Hidden attribute set, but this setting is not a function of its being a rootkit file; instead, it's just an extra precaution that malware writers (in general) often use to make their files less detectable (using the ordinary means of hiding a file in Windows). You need to reverse that setting to facilitate removal of any file or folder — especially malicious ones.



To set Windows to show hidden files and folders, choose Tools ⇨ Folder Options ⇨ View in Windows Explorer. Then check Show Hidden Files and Folders and uncheck Hide Protected Operating System Files. *These settings should be reversed at the end of any clean-up, for safety reasons.*

You can also use a DOS command to reverse the hidden attribute on a specific file or folder, by issuing the following at the command line:

```
Attrib -h -r <hiddenfile.xxx>
```

To issue this command, you must be in the folder that the file is located in, and you must substitute your own filename or folder for the generic filename we've specified within the brackets. The `-h` attribute makes the file unhidden, and the `-r` removes the read only attribute — so the file can be deleted,

Another approach you can take to determine whether a file is legitimate or infected is to take advantage of online scanning services. They allow you to upload a suspect file and test it for its threat potential. Here are a couple of scanners that do this job extremely well, along with sites where you can access them:


✔ **Virus Total:** www.virustotal.com/flash/index_en.html

✔ **Jotti:** <http://virusscan.jotti.org/>

The Virus Total and Jotti scanners are multi-threat scanners that submit the file you've uploaded to about 20 individual antivirus scanners, so you may get a cross-section of opinions in the final scan report. Virus scanners depend on databases that must be updated, and timing is crucial when a new threat is introduced. Because timing is of the essence, it's possible that some of the represented scanners may not have had their signatures updated to recognize a specific brand-new threat at the time of your testing. Therefore, using so many scanners tends to moderate this effect. Figure 9-9 depicts a set of scan results that Virus Total returned on an infected compacted folder. Notice how the scanners can give mixed results — and how the first three scanners, AntiVir, Avast, and AVG (all offered free to home users), were all able to correctly identify the threat.

Killing rootkits with IceSword

When process(es), autostart entries, services and drivers are confirmed to be of rootkit origin, IceSword enables you to deal with all of these items from within the IceSword program. IceSword provides both process and service control functions. Right-clicking a service listed in the Win32 Services display enables you to stop and disable it from the context menu. Right-clicking a process in the Process display enables you to terminate the process or even force-kill individual process threads.



The screenshot shows a browser window displaying the VirusTotal website. The main content is a table with the following columns: Antivirus, Version, Update, and Result. The table lists various antivirus engines and their scan results for a specific file.

Antivirus	Version	Update	Result
AntVir	6.35.0.13	06.18.2006	TR/Spy.KStaff.3
Authentium	4.93.8	06.16.2006	no virus found
Avast	4.7.844.0	06.15.2006	Win32:Zapchast-D
AVG	386	06.16.2006	LowZones.A
BKDefender	7.2	06.18.2006	Trojan.Bat.Zapchast.D
CAT-QuickHeal	8.00	06.17.2006	no virus found
ClamAV	devel-20060426	06.18.2006	Trojan.Zapchast-6
DrWeb	4.33	06.18.2006	Trojan.DownLoader.1844
eTrust-InoculateIT	23.72.42	06.18.2006	no virus found
eTrust-Vet	12.6.2259	06.16.2006	no virus found
Ewido	3.5	06.18.2006	Trojan.Zapchast
Fortinet	2.77.0.0	06.18.2006	W32/Winreg.A!tr
F-Prot	3.16f	06.17.2006	no virus found
IKarus	0.2.65.0	06.16.2006	no virus found
Kaspersky	4.0.2.24	06.18.2006	Trojan.BAT.Zapchast
McAfee	4786	06.16.2006	Downloader-QG
Microsoft	1.1441	06.18.2006	Trojan:WinREG/Lowzones.Q
NOD32v2	1.1606	06.17.2006	Win32/Adware.MediaTickets.downloader
Norman	5.90.21	06.16.2006	no virus found
Panda	9.0.0.4	06.18.2006	Adware/MediaTickets
Sophos	4.06.0	06.18.2006	Troj/WinREG-B
Symantec	8.0	06.18.2006	Trojan.LowZones
TheHacker	5.9.8.162	06.18.2006	no virus found
UNA	1.83	06.16.2006	no virus found
VBA32	3.11.0	06.18.2006	Trojan.BAT.Zapchast#32
VirusBuster	4.3.7.9	06.18.2006	BAT.Lowzones.Q

Figure 9-9:
Virus Total
scan
results.

After the rootkit has been disabled, clicking the File function opens a file-system browser (similar to Windows Explorer) so you can delete rootkit files and folders. The Registry function opens a Registry browser (similar to the Windows Registry Editor — Regedit) so rootkit autostarts and Registry entries may be removed. Deleting autostart entries is probably better left to a program like Autoruns, which deletes the startup for you and saves you from having to resort to the Registry editor. If you do choose to manually edit the Registry, be sure to back up a copy of the Registry (or the key you plan to edit) first. Any changes to the Registry are permanent and can only be reversed by using a backup copy of the Registry or a key. It's worth noting that these Registry entries represent nonfunctional orphans; the processes and services they refer to cease to exist after you've disabled and removed them. Therefore they represent no danger, and removing them is a straightforward housekeeping function.

After rootkit services and processes are stopped, you can choose File ⇨ Reboot to make any residual rootkit components visible and removable by traditional scanning programs. A good practice to follow is to reopen IceSword and confirm that no hidden items remain after a reboot.



Like BlackLight, IceSword can reverse DKOM modification of the Process List — which makes hidden processes visible. (Note, however, that it cannot detect processes that were hidden using PspCid Table modification.)

IceSword functions

IceSword provides a range of functions. Many of these functions list hidden or rootkit items in red, as indicated by the function descriptions listed below:

- ✔ **Process:** This option displays a list of active processes in the display hidden processes are indicated in red and are most likely of rootkit origin; other entries are listed in black.
- ✔ **Port:** This option lists active TCP and UDP endpoints along with information regarding these endpoints such as the state (listening), Process ID (PID) of the process that's maintaining this port, and the image data (path) of the program maintaining this open port. Programs such as Netstat can't see these rootkit ports.
- ✔ **Kernel Module:** This lists kernel-mode drivers. Color coding isn't used in the kernel module display. There are no red entries in this list to distinguish hidden rootkit drivers. If a rootkit driver is present, it's displayed — but it's listed in black.
- ✔ **Startup:** This lists autostart locations and any entries programs on your system may have inserted at these locations. The locations listed are similar to the ones flagged by MSConfig. This isn't a very exhaustive list; Autoruns and GMER can be used for a much more detailed listing of nearly every autostart location. However, one critical autostart location IceSword does show is the `AppInit_Dlls` key — a location commonly exploited by user-mode rootkits and general malware.
- ✔ **Win32 Services:** Lists all installed services and their status (running, disabled, stopped). Right-clicking a service brings up a context menu that gives you the option to Stop, Start, Pause, or Resume as a Service, depending on its present state. It also allows you to modify its Startup type (Manual, Automatic, Disabled). Hidden services are listed in red — and are almost always of rootkit origin.
- ✔ **SPI:** Displays Winsock Layered Service Providers, both good and bad. The Winsock chain is targeted by threats that seek to control or redirect your browser without your consent.
- ✔ **BHO:** Lists Browser Helper Objects, both good and bad. Legitimate programs (like helpful toolbars) can install BHOs, but threats can also install BHOs to control and modify your browser behavior.
- ✔ **SSDT:** Displays drivers with entries in the System Service Descriptor Table. Red entries have hooked the SSDT. Because many legitimate programs hook the SSDT, this doesn't provide a definitive diagnosis of rootkit behavior. All red results should be cross-correlated with those displayed by using the Win32Services, Process, and Kernel Modules functions to make a definitive rootkit diagnosis.

- ✔ **Message Hooks:** Displays programs that use active message hooking. There is a mix of both good and bad on this list. You can use the Message Hooks function to detect keyloggers that use DLL injection to capture keyboard input (we talk about that in more detail a bit later).
- ✔ **Log/Process Thread Creation:** Displays a log of process thread and activity with a date and time stamp. Suspect items are highlighted in red, but they don't provide a definitive diagnosis of rootkit behavior. Processes launched through another running process are colored (blue, black, green = safe, red = investigate). Legitimate system processes may appear in red (especially `svchost.exe`) because of the DLLs they load. Log/Process Thread Creation can be used in combination with the Reboot and Monitor function (on the Menu under File — Options). Choosing Reboot and Monitor initiates an immediate reboot and logs all process activity during startup. After rebooting, when you reopen IceSword and choose Log Process and Thread Creation, you can view process activity from system startup to the present.
- ✔ **Log Process Termination:** Creates a time log of terminated processes.
- ✔ **Registry:** Opens IceSword's Registry Editor so you can delete rootkit entries.

If you're familiar with the Registry, you can inspect the services key: `HKLM\SYSTEM\CurrentControlSet\Services\` to see whether any rootkit services have been added that aren't visible in the Win32 Services or Kernel Module functions. The rootkit `pe386` causes this apparent anomaly — because its driver hides in the ADS portion of the Windows `system32` folder, it's not visible in IceSword's Win32 Services or Kernel Module functions — but its service autostart is clearly visible in IceSword's Registry function. (Gotcha!)
- ✔ **File:** Opens IceSword's file-system browser so you can delete rootkit files and folders.



UnHackMe

Gratis UnHackMe is definitely the easiest and friendliest rootkit-detection program to use — especially if you're not all that computer-knowledgeable or you are a rootkit newbie. Its operation is extremely simple, and it provides removal as well as detection capability, which makes it an ideal beginner's choice. UnHackMe can effectively detect and eliminate HackerDefender, as well as other rootkits.



A full-featured evaluation version of UnHackMe is provided on your DART CD.

When you first open the UnHackMe program, rootkit detection is initiated by clicking the Check me Now button. UnHackMe then displays hidden entries and triggers an alert if the hidden files belong to a "Suspicious Trojan Class." This makes identification of rootkit components a no-brainer.

UnHackMe separates file and folder entries from Registry entries by displaying them in two separate panes. You can select AutoRemoval — or you can play a role in the removal process by indicating which actions you want to perform on the items you select from those displayed. The actions that can be applied to the items you've selected (by using Function buttons), are:

- ✔ Disable Autorun (deletes a program's autostart entry in the Registry)
- ✔ Delete File at next Reboot (used to delete files that are "in use")
- ✔ Stop Service (stops a process or driver that runs as a service)
- ✔ Delete Key (deletes a Registry key)

UnHackMe prompts you to confirm any deletion actions you've decided to take. After that, you can click Exit and reboot to delete any files you've marked for deletion (by selecting Delete File at Next Reboot). After you reboot, the rootkit files you marked for deletion are eliminated.



A demo — plus a very good help function — can be found on the Greatis Software Web site at www.greatis.com/unhackme/. Another excellent step-by-step screenshot tutorial, provided by the University of Minnesota, can be found at http://safecomputing.umn.edu/guides/scan_unhackme.html.

Malicious Software Removal Tool

The Microsoft Malicious Software Removal Tool (MSRT) automatically and silently scans your computer for a variety of prevalent threats — including worms, backdoor trojans, and commonly found rootkits — when you download Microsoft Windows Updates. The MSRT provides both detection and removal of any threats found. Most users are totally unaware that this process is happening — and that they've passed an important milestone when they install updates without triggering a threat alert. A new version of the MSRT is released with Windows Updates on the second Tuesday of every month.

To see a list of the threats that the MSRT currently scans for, please refer to this link:

```
www.microsoft.com/security/malwareremove/families.aspx
```

Why wait for another month to pass — you can perform an online MSRT scan more frequently by visiting the MSRT online scanner at this link:

```
www.microsoft.com/security/malwareremove/default.aspx#run
```



You must use a browser that supports ActiveX, such as Internet Explorer, to perform the MSRT online scan.

AntiHookExec

A more advanced tool than has been discussed so far in this chapter — and one of Nancy’s personal favorites — is AntiHookExec. AntiHookExec allows you to visualize and remove hidden rootkit components, using traditional system-analysis tools that you’re already comfortable using (and more apt to run on a regular basis). AntiHookExec when combined with these tools can provide early detection or an early warning of rootkit activity. This is valuable because it can help you discover a rootkit before it becomes deeply entrenched on your system.

Using AntiHookExec enables you to see the results in a matter of minutes — or even seconds — depending on which tool it’s combined with. It provides an important role as a rootkit screening tool to spot abnormal rootkit entries that might otherwise have gone unnoticed for a considerable amount of time. If rootkit components are detected, then other dedicated rootkit programs may be called into action as well, to confirm and eliminate the threat.

As its name implies, AntiHookExec works by restoring the data structures that have been hooked by rootkits, to their original unhooked state. When these structures are restored, traditional system-analysis tools (such as HijackThis, Autoruns, Regedit, Task Manager, and Process Explorer) can locate and remove rootkit components. AntiHookExec even creates a log of all the hooks that were restored for informational purposes. AntiHookExec restores user-level hooks (IAT, EAT) — including inline function hooks. It does not restore kernel hooks, so you’ll have to augment its capabilities by running a dedicated anti-rootkit program to detect kernel-mode rootkits.



AntiHookExec is not meant to be the only anti-rootkit tool in your arsenal — but if you launch all your diagnostic programs from the run line in combination with AntiHookExec, you’ll be sure to catch any user-level rootkit activity in the act. Nancy launches all her system-analysis tools this way.

After we show you how to install the program, we give you some concrete examples that demonstrate exactly how AntiHookExec works.

AntiHookExec is available as a free download from the SIG² Web site:

www.security.org.sg/code/antihookexec.html

Installing AntiHookExec

There are two methods of installing AntiHookExec so it works properly with other programs. Use only Method 1 or Method 2 — not both.

Method 1

The unzipped program file called `AntiHookExec.exe` must be moved to the `C:\Windows\System32` folder

Method 2

This alternative updates the `path` environment variable so it recognizes the `AntiHookExec` directory, as a system-wide variable.

You can access the `Path` variable by right-clicking `My Computer` and selecting `Properties` ⇨ `Advanced` ⇨ `Environment Variables`. Then double-click `Path`, and in the dialog box that opens, append the `AntiHookExec.exe` folder location to the end of the path value.

For example, if `AntiHookExec.exe` is in the `C:\Program Files\AntiHookExec\` folder, you would append the following command to the end of the path, including the semicolon:

```
;C:\Program Files\AntiHookExec\
```

Using AntiHookExec with other tools

If you're familiar with Windows processes and services, and have a fair knowledge of system-analysis tools such as `HijackThis`, `Process Explorer`, `Autoruns`, and `Port Explorer`, this utility can greatly ease interpretation of results. This technique can also help you recognize a rootkit early in its life cycle — before it has inflicted too much harm.

If you think you're infected by a stealth threat that isn't visible to conventional analysis tools — or if you're experiencing problems even though all your scans are returning clean results — then by all means run several of the tools listed here (the way we suggest in the following section), and correlate their findings.

The following sections show you a few examples of how to use `AntiHookExec` with traditional system-analysis tools so you can get the flavor of its power. In all these examples, we're using tools that would not ordinarily be capable of identifying rootkit components, had they been launched without using `AntiHookExec`. Our “test” rootkit used with all the tools (`HijackThis`, `Autoruns`, `Process Explorer`) in our examples is the `HackerDefender` rootkit (a.k.a. `HxDef`) — `HackerDefender` installs a service that launches a process called `hxdef100.exe`. It also installs a driver called `HxDefdrv.sys`. `HackerDefender` is a hybrid rootkit (part user-mode and part kernel-mode — hence

the presence of the driver component). HxDf's process hides its resources (files, Registry keys). The HxDf driver is used to give HackerDefender backdoor functionality. It does this by intercepting inbound traffic and hijacking a port that's being used by another application. By hijacking an existing port, HackerDefender's backdoor manages to operate without firewall interference and camouflages its activities. This mechanism makes it unnecessary for HackerDefender to hide its port using rootkit techniques.

Although identifying and stopping the HxDf service (process) is enough to uncloak HackerDefender, ideally the best tools to use are those that are capable of identifying both the rootkit driver and process — so the driver file can be deleted, too. After the rootkit components are identified they must be removed, and we show you different manual methods of doing that. Because many rootkit detectors identify rootkits but offer no means of removing them (Rootkit Revealer, SVV), these same removal techniques can be applied when using those types of detectors as well.

Using AntiHookExec with HijackThis

The drawback to using the HJT scan alone is that the driver file isn't visible in the HJT log but it *is* visible in the HJT Startup List under the Enumerating Windows NT/2000/XP services category. Still, performing a routine HijackThis scan reveals the HxDf service called HXD Service 100. When the process associated with this service is stopped, it effectively uncloaks the rootkit, so we can locate its files on disk and remove them.

To launch HijackThis through AntiHookExec from the run line click Start, choose Run, type **AntiHookExec "C:\Program Files\Hijackthis\hijackthis"** into the Open: field, and then press Enter or click OK. (This assumes HijackThis.exe is located in the C:\Program Files\Hijackthis folder.) The quotes are required because of the space in the HijackThis path. Naturally, adjust the path accordingly to reflect the location of HijackThis.exe on your system.

HijackThis displays the running HackerDefender service called HXD Service 100 under the 023 Service entries. The HackerDefender service is disabled and stopped by checking the 023 Entry and clicking Fix Checked.

Because HJT provides the path and filename of the service executable file, it can be removed in Windows Explorer after the service is stopped.

As Figure 9-10 illustrates, AntiHookExec used with HijackThis (HJT) can display the (normally hidden) rootkit service.

```

hijackthis.log - Notepad
File Edit Format View Help
LogFile of HijackThis v1.99.1
Scan saved at 12:05:25 AM, on 6/18/2006
Platform: windows xp sp2 (winnt 5.01.2600)
MSIE: Internet Explorer v6.00 SP2 (6.00.2900.2180)

Running processes:
C:\WINDOWS\system32\smss.exe
C:\WINDOWS\system32\winlogon.exe
C:\WINDOWS\system32\services.exe
C:\WINDOWS\system32\lsass.exe
C:\WINDOWS\system32\svchost.exe
C:\WINDOWS\system32\svchost.exe
C:\WINDOWS\explorer.exe
C:\WINDOWS\system32\spoolsv.exe
C:\Program Files\VMware\VMware Tools\VMwareService.exe
C:\Program Files\VMware\VMware Tools\VMwareTray.exe
C:\Program Files\VMware\VMware Tools\VMwareUser.exe
C:\Program Files\Sandboxie\Control.exe
C:\Documents and Settings\swat\My Documents\hxdef100r\hxdef100.exe ←
C:\WINDOWS\system32\cmd.exe
C:\Tools\HijackThis\HijackThis.exe

04 - HKLM\...\Run: [VMware Tools] C:\Program Files\VMware\VMware Tools\VMwareTray.exe
04 - HKLM\...\Run: [VMware User Process] C:\Program Files\VMware\VMware Tools\VMwareUser.exe
04 - HKCU\...\Run: [SandboxieControl] C:\Program Files\Sandboxie\Control.exe
09 - Extra button: Messenger - {FB5F1910-F110-11d2-BB9E-00C04F795683} - C:\Program
Files\Messenger\msmsgs.exe
09 - Extra Tools\menuitem: windows Messenger - {FB5F1910-F110-11d2-BB9E-00C04F795683} -
C:\Program Files\Messenger\msmsgs.exe
020 - Service: hxp service 100 (HackerDefender100) - unknown owner - C:\Documents and ←
Settings\swat\My Documents\hxdef100r\hxdef100.exe
023 - Service: VMware Tools Service (VMTools) - VMware, Inc. - C:\Program Files\VMware\VMware
Tools\VMwareService.exe

```

Figure 9-10:
AntiHook
Exec used
with
HijackThis.

Two HijackThis functions available in the Open the Misc Tools section called Delete an NT Service and Delete a file on Reboot can remove both the rootkit service autostart and its associated file, respectively. The final result is that we can identify and delete the HackerDefender rootkit service by using HijackThis alone, when HijackThis would normally not even see the rootkit entries. Because the HackerDefender service loads its driver, deleting the service is enough to disarm the rootkit (if the service isn't running, it can't load the driver). If we created a startup list, we could identify the rootkit driver — and delete that driver — using the same HijackThis option (Delete File on Reboot) that we used to delete the HackerDefender process's executable file. All this capability is made possible by using HijackThis in combination with AntiHookExec. If we had not run them together, none of the HackerDefender entries would have been visible.



If you're interested in learning how to create a HJT Startup list, please refer to the directions in this link:

http://wiki.castlecops.com/The_HijackThis_Startup_List

Using AntiHookExec with Autoruns

Here's where we show you how to use AntiHookExec in combination with Autoruns to target rootkit autostarts and the processes (files) launched by them. In our example, AntiHookExec allows us to see the Registry autostarts for HackerDefender and the filenames and paths (images) associated with those entries.

To launch Autoruns through AntiHookExec, click Start, choose Run, type **AntiHookExec “C:\Program Files\autoruns\autoruns.exe”** into the Open: field, and then hit Enter or Click OK.

Bingo! When Nancy ran this on a test computer using the default Everything Tab in Autoruns to display *all* startup categories, Autoruns displayed both the HxDf Driver (`hxdefdrv.sys`), and the HxDf service (`hxdef100.exe`) autostarts under the Services Registry key:

```
HKLM \SYSTEM\CurrentControlSet\Services\
```

Both items are listed under the Autorun Entry column as HackerDefender. The image path column indicates the location of the rootkit files on disk. Figure 9-11 illustrates the rootkit service startup visible in Autoruns with the Services tab selected.




Figure 9-11:
AntiHook
Exec used
with
Autoruns.

Now that you know the path of both `hxdef100.exe` and `hxdefdrv.sys`, you can delete them manually in Windows Explorer after the service is stopped. To stop the service, you first have to remove the service autostart entry, and then reboot. Removing the autostart entry in Autoruns is very easy — just right-click it, and choose delete. Now, when you reboot your computer, the process that the service started won't be running (because the autostart was removed), and you can then delete the service and driver files on disk.



Alternatively, you can use Process Explorer and Autoruns *together* to stop the `hxdef100.exe` process (to eliminate the reboot). To do that, launch Process Explorer in combination with AntiHookExec using the command specified in the section that follows (leave Autoruns open). Next, right-click the service autostart entry in Autoruns, and choose Process Explorer from the context menu. This will immediately open the Properties Dialog associated with the service entry in Process Explorer. Click the Stop button, and the service will be stopped! Because Process Explorer and Autoruns were both developed by Sysinternals, they work together as a duo to facilitate removal of unwanted processes and services. Figure 9-12 shows Process Explorer's Process Properties dialog box for `hxdef100.exe` with the Services tab selected.




Figure 9-12:
The
Properties
dialog box in
Process
Explorer,
showing
`hxdef100.exe`.

Regardless of the method you choose to stop the service (Autoruns or Process Explorer), the following HxDef files would need to be deleted using Windows Explorer (using the locations that were identified using the Everything tab in Autoruns):

- ✓ **Driver Image Path:** `C:\Documents and Settings\swat\My Documents\hxdef\hxdefdrv.sys`
- ✓ **Service Image Path:** `C:\Documents and Settings\swat\My documents\hxdef\hxdef100.exe`

With the rootkit service no longer running, the autostarts become visible in Autoruns when you launch it normally without AntiHookExec. When in Autoruns, you can delete the rootkit autostarts by highlighting each autostart entry and then choosing Entry ⇨ Delete. Alternatively, you can highlight each autostart entry and choose Delete from the context menu.

Using AntiHookExec with Process Explorer

To launch Process Explorer through AntiHookExec, click Start and choose Run, and then type **AntiHookExec “C:\Program Files\Process Explorer\procexp”** into the Open: field (remember to include the quotation marks). Then hit Enter or click OK.

Remember to substitute the correct path for Process Explorer on your system.

Process Explorer displays the running service indicated by the `hxDef100.exe` entry in the Process Tree. Right-clicking `hxDef100.exe` and selecting Properties from the context menu brings up a Properties dialog box for the `hxDef100.exe` process. The Properties dialog box indicates the location of the file on disk so you can manually delete it in Windows Explorer, after the rootkit service is stopped. To stop the rootkit service, click the Services tab in the `hxDef100.exe` Properties dialog box and click Stop. (This method duplicates our shortcut approach used in the previous Autoruns example.)

Figure 9-13 illustrates how AntiHookExec can be used with Process Explorer to detect HxDef’s rootkit process (service) `hxdef100.exe` in Process Explorer’s Process Tree.




Figure 9-13:
AntiHook
Exec used
with
Process
Explorer.



What you can't see in the grayscale image shown in Figure 9-13 is that Process Explorer differentiates running services from running processes by highlighting services in pink and active processes in lilac.

Next, to see whether any rootkit drivers are installed, return to the Process Tree View in the Process Explorer Main Display; and then follow these steps:

1. Highlight the System Process (System) in the Upper Pane.

The Lower Pane gradually updates to display all kernel-mode drivers on your system.

2. Choose Options → Verify Image Signatures to have Process Explorer verify all drivers for Digital Signatures. (There may be a delay while signature information is obtained.)

This makes the HxDf Driver easy to spot in the Lower Pane, since nearly all the Drivers entries listed bear a *Verified* label.

3. Locate the HxDf Driver in the Lower Pane (`hxdefdrv.sys`) and right-click the entry.

A Driver Properties dialog box appears.

4. Select the Image Tab to display the location of `hxdefdrv.sys` on your hard drive.

Because the `hxdef100.exe` service has already been stopped, the driver has been unloaded and it becomes visible for deletion within Windows Explorer.

5. Use the location you found in Step 4 to find and delete the file in Windows Explorer.

VICE

VICE (Virtual Intruder Capture Engine) is a rootkit-detection tool written by Jamie Butler, the same person who developed the FU (proof-of-concept) rootkit. VICE may be downloaded here:

www.Rootkit.com/vault/fuzen_op/vice.zip

VICE can detect both user and kernel-mode rootkits that utilize hooking techniques to hide. It doesn't do a cross-view comparison (as Rootkit Revealer does); rather, it relies on heuristics (rootkit-type behavior) analysis to determine the presence of a rootkit.

VICE analyzes the data structures in each user program's address space (IAT, EAT) to determine whether any API hooking has occurred. It also detects user- and kernel-mode inline function hooking of the APIs and kernel services referenced by the IAT and SSDT, respectively, that have had code added to facilitate the execution of replacement rootkit functions. VICE detects hooking of global (kernel) data structures vulnerable to alteration by kernel-mode rootkits — specifically the SSDT and the IRP (I/O Request Packet) Tables of individual drivers.



The operating system and user applications use I/O Request Packets to communicate with drivers when they need to execute a specific function that a driver supplies. Each kernel driver has an IRP Function Table (I/O Request Table) that contains pointers to functions to be executed in response to the IRP requests. VICE checks to see whether the IRP Table function pointers have been replaced to point to alternate functions that lie outside the driver's address space.

Like most rootkit-detection programs, VICE relies on heuristics or behavioral analysis to detect rootkits. The drawback of programs that rely on heuristics is that they can lead to false positives, because they often don't take into account legitimate programs that hook kernel data structures to accomplish their goals. For example, the very popular firewall program ZoneAlarm uses SSDT hooking to filter Internet traffic. However, malware applications may maliciously hook the SSDT to intercept kernel services as well. Because both legitimate and malicious programs hook kernel data structures, care must be taken to scrutinize the VICE scan results to rule out false positives. VICE does flag many valid Windows DLLs in its scan results (`sfc.dll`, `adsldpc.dll`, `setupapi.dll`, `shim.dll`).

What's nice about VICE (we just had to say that), besides its excellent detection capabilities, is that it specifies what is being hooked and identifies the path of the file responsible for doing the hooking. If an anti-rootkit program doesn't do that the scan results can be very difficult to interpret. Luckily, VICE does track back and identify the rooted files — not only does this simplify the process of distinguishing false positives from the real thing, but it also enables you to readily locate and remove rootkit files on your hard drive.



The Microsoft .NET Framework is required to support the VICE graphical user interface. It may be downloaded for free from Microsoft.

System Virginty Verifier (SVV)

System Virginty Verifier (or SVV) was developed Joanna Rutkowska, a leading expert in stealth technology and owner of invisiblethings.org. SVV may be downloaded at the invisiblethings.org Web site.

SVV uses a heuristic approach similar to that of VICE. It checks the user-program and kernel data structures that are prone to rootkit modification (such as the IAT, EAT, SSDT, and IRP tables). Like VICE, it also checks for inline function hooking of APIs in user-level DLLs and the kernel services referenced by the SSDT. SSV does memory-verification checking and also attempts to reduce false positives by distinguishing *innocent hooking*, used by legitimate security program drivers, in its scan results.



To make SSV easier to access, it's best to install it to your computer's root directory (which is `C:\` on most systems).

SVV has to be run from the command console as follows:

1. **Click Start and select Run, type `cmd` into the Open: field, and click OK or press Enter.**

The command prompt appears.

2. **Type `cd\` to get to the root directory (usually `C:\`).**
3. **Type `svv check /a /m` and then hit Enter.**

This command tells SVV to check all modules and show all details about any detected modifications. SVV will issue a warning, and ask whether you want to continue.

4. **Type `yes` and hit Enter to continue.**

SVV responds by saying `Warming up`, which indicates it's actively scanning. When SSV finishes, it post its findings.

The SSV scan report lists suspect entries and rates them with a color-coded system of six threat levels, ranging from BLUE to DEEPRED. BLUE indicates no suspect activity; DEEPRED indicates probable infection.

Of the three sample SVV logs illustrated here, Log 1 represents a completely clean scan; and Logs 2 and 3 indicate probable infection. Here's an example of a clean SSV log:

```
SYSTEM INFECTION LEVEL: 0
0 - BLUE
1 - GREEN
2 - YELLOW
3 - ORANGE
4 - RED
5 - DEEPRED
```

This is an example of an infected SSV log showing API hooking:

```
E:\winapps\svv>svv check
ntoskrnl.exe (804d7000 - 806eb400)... innocent hooking (verdict = 2).
NDIS.SYS (f765c000 - f7689000)... innocent hooking (verdict = 2).
kernel32.dll (7c800000 - 7c8f4000)... suspected! (verdict = 5).
WS2_32.dll (71ab0000 - 71ac7000)... suspected! (verdict = 5).
USER32.dll (77d40000 - 77dd0000)... suspected! (verdict = 5).

SYSTEM INFECTION LEVEL: 5
0 - BLUE
1 - GREEN
2 - YELLOW
3 - ORANGE
4 - RED
5 - DEEPRED
SUSPECTED modifications detected. System is probably infected!
```

The first two entries indicate hooking of `ntoskrnl.exe` and `NDIS.SYS` by legitimate programs and are given a verdict of 2 (innocent hooking)

The last three entries are suspect and are marked with a (`verdict=5`) to indicate that. They represent hooked APIs in user-level DLLs.

This is an infected log that shows an inline function hook in `Kernel32.dll` (notice the `JMP` code label on the fourth line shown in bold):

```
kernel32.dll (7c800000 - 7c8f4000)... suspected! (verdict = 5).
module kernel32.dll [0x7c800000 - 0x7c8f4000]:
0x7c801af1 (section .text) [LoadLibraryExW()+0] 6 byte(s):
JMPing code (jmp to: 0x5f05001e) <== inline hook
address 0x5f05001e DOES NOT belong to ANY MODULE!
file :6a 34 68 58 e0 80
memory :ff 25 1e 00 05 5f
verdict = 5
0x7c80abf3 (section .text) [FreeLibrary()+15] 4 byte(s): SUSPECTED code
modification file :dc ff ff ff memory :45 54 7f e2
verdict = 5
module kernel32.dll: end of details
SYSTEM INFECTION LEVEL: 5
0 - BLUE
1 - GREEN
2 - YELLOW
3 - ORANGE
4 - RED
--> 5 - DEEPRED
SUSPECTED modifications detected. System is probably infected!
```

SVV can attempt to fix the modifications that it noted in its scan results. To do that, you have to type **svv fix** and hit Enter. In response, SVV issues a warning and asks you if you're sure you want to continue. Again, you must type in **yes** and then hit Enter to continue. SVV then attempts to fix the items that were flagged as RED or DEEPRED in its scan results.



If you find you're getting false-positive detections in your SVV scan results, you should completely disconnect from the Internet, and then turn off all active-protection programs — but leave your firewall engaged. This will eliminate interference caused by security-program drivers. Then repeat the scan.

Strider GhostBuster

Strider GhostBuster is a Microsoft-conceived rootkit detector that offers two different scanning options. Each scanning option uses a cross-difference comparison technique, as follows:

- ✓ **Inside-the-box GhostBuster:** Features a low-level-to-high-level Windows API cross-difference comparison scan.
- ✓ **WinPE GhostBuster:** Features an internal to external cross-difference comparison scan.

Unfortunately, GhostBuster is still in the experimental or prototype phase; it's currently unavailable for download from Microsoft at the time of this writing. GhostBuster sounds like a very promising rootkit detector; hopefully it (or a very similar program) will become available to users of Microsoft operating systems in the near future.

Information on the Microsoft Strider GhostBuster detection utility may be found here:

```
http://research.microsoft.com/Rootkit/
```

Inside-the-box GhostBuster

GhostBuster conducts its inside-the-box (internal) system scan from within the host system using low-level disk reads, and a high-level scan using Windows APIs similar to Rootkit Revealer. Unlike Rootkit Revealer, it has an added bonus — GhostBuster also analyzes kernel data structures, using similar low-level and high-level techniques to detect hidden processes. By doing a cross-difference comparison of the results generated by these two scanning methods, GhostBuster can detect discrepancies possibly caused by rootkit modification. Except for the added hidden process-detection feature, this method is analogous to the one employed by Rootkit Revealer.

WinPE GhostBuster

WinPE GhostBuster performs an internal-external comparison scan of the host computer. First, a scan is performed from within the host computer and the results are saved on the hard drive of the host system. Next, the host computer is rebooted to an external Win32 PE CD-ROM and an identical external scan of the host computer is performed. By booting to an external medium, an uncompromised view of the host computer is obtained in the external scan results. A difference comparison of these inside and outside views is then made from the external CD-ROM. Any discrepancies can be attributed to rootkit or stealth programs resident on the host computer.

Rootkitty

Although GhostBuster isn't available for release yet, Rootkitty has automated the steps recommended by Microsoft on the Strider GhostBuster Web site, to perform an inside-the-box and outside-the-box cross-difference file comparison scan. Rootkitty is a hidden program-detection tool included in the UBCD4Win CD build, and it's also available for download on the UBCD4Win Forum:

```
www.ubcd4win.com/forum/index.php?showforum=48
```

Rootkitty uses the procedure advocated by Microsoft's Strider GhostBuster Web page to detect *ghostware* (a general term Microsoft uses to describe unwanted and uninvited stealth programs). Under the Tools section, Microsoft recommends running some simple scans using DOS directory commands from inside and outside the suspect system. The exact procedure they recommend can be found at the Microsoft's Strider GhostBuster Web page at

```
http://research.microsoft.com/rootkit/
```

Here's a summary of the simple steps you can take — or let Rootkitty perform automatically — to detect some of today's ghostware. In essence, two scans are conducted from within the potentially infected host OS and the results are saved. Then the same scans are performed from outside the host system and again the results are saved. Comparing the internal and external scan results discloses stealth files. The sequence looks like this:

1. The directory command detects all hidden files on the hard drive in brief format, yielding just enough information to facilitate a simpler difference comparison. The first command is

```
dir /s /b /ah.
```

The results are saved to a file.

2. The second command runs exactly the same scan, this time excluding hidden files from the output. The second command is

```
dir /s /b /a-h
```

The results are saved to a file.

3. These two scans together represent all files (including hidden ones) on the host computer's hard drive, using an inside-the-box view.
4. The identical scans are repeated (using the same commands) on the hard drive, but they're conducted from a clean bootable CD-ROM instead.

These two external scans together, represent all files (including hidden ones) on the hard drive from an outside-the-box view.

5. The inside-the-box and outside-the-box views are compared, using a program called WinDiff (run from the bootable CD-ROM), and the differences are recorded in a scan report.

These differences represent any files that were invisible from within the infected computer but visible from the outside — and that's your ghostware (stealth malware).



False positives may appear in the difference results. A drawback to this technique is that it can't identify malware that hides in unusual hiding places to avoid detection — for example, ADS (alternate data streams) of files or folders on NTFS systems, the BIOS (Basic Input/Output System), track 0 of the MBR (Master Boot Record), video-card EEPROM (Electrically Erasable Programmable Read-Only Memory), or bad disk sectors.

Rootkitty has incorporated the procedure shown here into a simple program that scans your system both from the inside and outside and then does a cross-difference comparison between the scans. (This relieves interested parties from having to implement the steps advocated by Microsoft.) Rootkitty does two essential comparisons: An inside-to-outside file-difference comparison on the infected target (using the procedure just described) and a cross-difference comparison of MD5 checksums for each file. (The “Checkpointing Utilities with Offline Hash Databases” section earlier in this chapter gives you a closer look at that process.)

RAIDE

We provide details of RAIDE in Chapter 7. Because it's only available to the public as a limited Beta release, we briefly review its approach here.

RAIDE is a very advanced anti-rootkit program that targets the most sophisticated techniques used by kernel-mode rootkits — namely DKOM. RAIDE also

detects virtual memory hooking. It was developed to counteract the new subversion techniques built into the laboratory-designed rootkits FUTO and Shadow Walker (see Chapter 12 for details). RAIDE can detect the hooking of kernel data structures (such as the SSDT) — and it can restore both the Process List and the Process and Thread List (PspCid Table) to reveal hidden processes. Shadow Walker, FUTO, and HackerDefender are all detectable by RAIDE. (In fact, RAIDE is the only program that *can* detect Shadow Walker at the time of this writing.)

Since RAIDE was developed, two anti-rootkit programs — DarkSpy Anti-Rootkit and GMER — were released. Both programs can detect FUTO's process- and driver-hiding techniques as well. To sum up, RAIDE can detect

- ✓ Processes and drivers hidden using DKOM
- ✓ SSDT hooking
- ✓ DKOM of tokens used to suppress the reporting of rootkit activity to Event Viewer
- ✓ Hooking of virtual memory by the memory-resident rootkit Shadow Walker

However, because RAIDE emphasizes advanced rootkit-detection techniques, it doesn't detect hidden files, folders, or Registry keys. It also doesn't detect the hooking of Driver IRP (I/O Request Packet) Tables or of the Interrupt Descriptor Table (IDT). It's therefore advisable — for complete coverage — to run RAIDE in combination with other rootkit programs that detect those items (such as DarkSpy, GMER, and IceSword).

DarkSpy

DarkSpy is a very powerful anti-rootkit program from China (developed by Mingyan Sun and Jianlei Shao) that's included on your DART CD. DarkSpy is similar to IceSword, in that it offers many different functions to detect and remove rootkits but it has some features that surpass IceSword's. In fact, many anti-rootkit program developers swear by DarkSpy — and that is the ultimate testament to its fine design.



Because DarkSpy scans your system at such a low level, it may interfere with security-program drivers, so it is best to disable active protection before using DarkSpy. When you run DarkSpy for the first time, it prompts you to choose between Normal Mode and Super Mode. Super Mode enables additional features that enhance DarkSpy's detection capabilities, but a reboot is required for Super Mode to become active. If you opt to run DarkSpy in Normal Mode, Super Mode will be enabled automatically the next time you reboot and launch DarkSpy.

DarkSpy's offers the following five functions to detect and remove rootkits, which you can access by choosing the following tabs:

- ✔ **Process:** This function displays all the running processes in the system. Hidden/rooted processes are shown in red.
- ✔ **Driver Module:** This function displays all the drivers loaded in the memory. Hidden rootkit drivers are shown in red.
- ✔ **File:** This function provides a file-system browser similar to Windows Explorer. The major advantage of this file-system browser is that it can show rootkit files because it performs low-level (raw) disk reads that bypass Windows APIs
- ✔ **Registry:** This function enables you to view and edit the Windows Registry. Because DarkSpy performs raw disk reads to analyze the Registry data, hidden entries inserted by rootkits are exposed.
- ✔ **Port:** This function shows all open ports (including hidden ones), and the local and remote IP addresses that are associated with them. You can use this feature to track backdoor trojans that maintain an open connection to a remote computer.

Detecting and removing rootkits with DarkSpy

Because DarkSpy lists only hidden processes in red, you should first click the Process tab to look for hidden rootkit processes. If you do find a red rootkit process, then you can terminate it by right-clicking it and choosing Kill. IceSword also offers a Force Kill option that you can use to end stubborn processes, but you should always try to use the Kill option first (to limit unanticipated side effects).

After you kill a rootkit process, you can use the File tab of DarkSpy to navigate to the folder where the rooted process is located, and then delete the file manually. Because DarkSpy's file function shows rooted files, it's not even necessary to kill a rootkit process to make it visible. However, because you can't delete a process if it's running, you'll still need to kill any red process before deleting it.

Because some kernel-mode rootkits have only a hidden driver and no processes, you should always be sure to use DarkSpy's Driver Module function to check for hidden drivers. DarkSpy really excels in its driver-detection capability; it can even identify drivers that are hidden using DKOM techniques (to remove the rootkit driver from the loaded drivers list and the object directory in kernel memory). The Driver Module function lists rootkit drivers in red, and the driver's file path is also noted. Once you know the driver's name, you'll need to remove the rootkit driver's Service key using DarkSpy's Registry function (see our "Analyzing the Registry" section to see how that's done). After the service key is deleted, the rootkit driver will be

unable to load after you reboot the computer. When the driver is deactivated in this way, all the resources it was hiding (files, folders, Registry entries, and ports) become visible, so you can manually delete them.

Using DarkSpy's Registry Analyzer

DarkSpy comes with a very advanced and multi-functional Registry Analyzer. The Windows Registry is composed of files on disk, and since DarkSpy can perform a low-level disk read of the Registry files, you can use it to dig out Registry entries created by rootkits. Even if a rootkit's Registry entries are invisible to Regedit (because they hook the Registry-related Windows APIs), they're visible when you use DarkSpy's Registry Analyzer.

Of all the methods that DarkSpy offers for analyzing the Registry, the one that's easiest to use (it's a lot like using Regedit) is the Online Analyze function. Here's the complete list of DarkSpy's Registry options:

- ✔ **Soft Save:** Saves the Registry branch or key to a specified binary file. DarkSpy uses some Windows APIs to fetch the data from the Registry files (hives) on disk.
- ✔ **Soft Restore:** Restores a specified Registry branch (hive) from data input from a previously saved binary file.
- ✔ **Raw Save:** Reads the selected Registry hive file directly from the disk, without using Windows APIs, and saves it to the file you specify in binary format. This file can be read later using DarkSpy's Offline Analyze function.
- ✔ **Boot Script:** DarkSpy can generate command-line scripts to fetch data from (or store data to) the selected Registry hive that you specify. These scripts can be used to analyze the Registry hives from a bootable CD. This feature is geared toward programmers, but the advantage is that it offers a totally uncompromised view of Registry data.
- ✔ **Offline Analyze:** DarkSpy loads a previously saved binary Registry file into its built-in Registry Analyzer. After that file is loaded, you can browse and edit the Registry keys or branches the same way you would if you were using Regedit. Thus, you can use a clean system to examine the Registry information obtained from an infected computer.
- ✔ **Online Analyze:** This is the function most closely resembles using Regedit to view and modify your Registry. When you choose this option, DarkSpy loads the Registry branch (*hive*) that you specify into its built-in Registry Analyzer using low-level (*raw*) reads. This technique enables you to see rootkit entries that are normally invisible, when you browse and edit the Registry.

Analyzing the Registry

Because most rootkits hide their Registry entries by hooking the Registry-related Windows APIs, the Windows Registry Editor (`Regedit.exe`) is of no help when you're trying to locate and remove cloaked Registry entries. However, you can use the Registry features provided by DarkSpy to perform a thorough and foolproof Registry analysis. Here's an example of what you can uncover; most rootkits install a service that is registered under these Registry branches:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services
HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Services
HKEY_LOCAL_MACHINE\SYSTEM\ControlSet002\Services
```

Here `CurrentControlSet` defines the service information that your computer uses to implement services during its normal operation — as represented by this Registry key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\
```

The information in the `CurrentControlSet` is derived from (and an exact duplicate of) one of these service keys:

```
HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Services\
HKEY_LOCAL_MACHINE\SYSTEM\ControlSet002\Services\
```



Windows maintains more than one control set for system-recovery purposes. If the `CurrentControlSet` results in a boot failure, often you can reboot the system successfully by choosing the Last Known Good Configuration from the Advanced Options Menu at startup. When Windows starts, it inspects the following Registry key to determine which `ControlSet` is going to be the current one:

```
HKEY_LOCAL_MACHINE\SYSTEM\Select
```

When you use DarkSpy's Registry function, DarkSpy notes which `ControlSet` is being used as the `CurrentControlSet` and shows that information at the bottom of its main display as follows:

```
CurrentControlSet: ControlSet001
```

Here the `CurrentControlSet` was derived from the information in `ControlSet001`. You can use that information to specify the control set you want to edit or view when you use DarkSpy's Registry analyzer. (DarkSpy gives you the option of editing only `ControlSet001` and `ControlSet002`.)

Here we show DarkSpy's offline and online Registry features in action in two examples: The first shows how you can use DarkSpy's Raw Save and Offline Analyze functions together; the second shows Online Analyze usage.

Example 1: Using Raw Save and Offline Analyze

First, we select the ControlSet that DarkSpy has indicated is the CurrentControlSet, choose the Raw Save button and provide a filename to save that Registry branch (hive) to on the hard drive. Next, we retrieve and analyze the service Registry information that was just saved, using DarkSpy's built-in Registry analyzer. To do that, choose Offline Analyze and provide the path of the saved file. DarkSpy opens the saved ControlSet hive that you've specified. From here, you can navigate through the services defined by this branch to find any suspicious or unknown Registry keys. Normally you would use this function to analyze portions of the Registry that were saved on an infected computer and then *ported* (transferred) to a clean system for analysis.

Example 2: Using Online Analyze

You can analyze the Registry hive online (without saving and loading the hive). Simply select the Registry hive you want to analyze and click Online Analyze. DarkSpy performs a low-level read of that hive and loads it into its Registry Analyzer. Then you can browse and edit that Registry information in the same way you would if you were using Regedit.



Keep in mind that editing the Registry is a very delicate operation, and we only recommend using DarkSpy's Registry feature if you're an advanced user, comfortable and familiar with editing the Registry. If you delete or change any legitimate key, then your computer may become unusable — so even if you're a Registry pro, you should still back up the Registry (or at least the key you're working on) before making any changes. (Chapter 3 discussed Registry backup basics.) DarkSpy provides an extremely strong system for Registry detection, but the tradeoff is that you must have a similarly strong knowledge to use its Registry features.

Using DarkSpy to remove difficult rootkits

DarkSpy's strength is that it can detect the most difficult DKOM and inline kernel-hooking techniques that abound today. DarkSpy can detect not only FUTO's hidden process (DKOM of both the Process List and PspCid Table) but also files and folders that are hidden by inline hooking system services in the kernel (`ntoskrnl`). DarkSpy also allows you to edit Registry entries that have had their access blocked by kernel inline function hooking. (Any applications that rely on system calls for Registry access are prevented from touching these rooted entries.)

To show off DarkSpy's strong detection features, the DarkSpy authors created their own demonstration rootkit called BadRKDemo. BadRKDemo uses the following rootkit subversion techniques:


- ✓ **Hides its driver using DKOM techniques:** As described in Chapter 7, DKOM modifies structures and objects that the kernel maintains in memory to keep track of its many activities. By manipulating these objects, a rootkit can distort the OS's view of what's really present on the system. BadRKDemo avoids detection by removing itself from the list of loaded drivers (`PsLoadedModuleList`) — and from the resources listed in the *object directory* (a structure that the kernel uses to keep track of open objects). By modifying *both* of these data structures, the DarkSpy demo is able to fool most rootkit detectors.
- ✓ **Blocks modification of its service Registry key:** By inline hooking a system service called `ObOpenObjectByName`, BadRKDemo can block detectors and Registry editors that use system calls to modify the Registry.
- ✓ **Achieves functionality solely through its driver:** BadRKDemo has no active process. It uses a process to load its driver, but then that process exits; only its driver and service Registry entry (autostart) remain.

These techniques allow BadRKDemo to hide from the System Service Controller and thereby prevent Registry access. Most rootkit detectors can't detect BadRKDemo's driver and are unable to remove its Registry autostart (though they are able to see it). Figure 9-14 shows how you can use DarkSpy's Driver Module Function to detect BadRKDemo's driver. Keep in mind that BadRKDemo hides its driver by modifying the driver list and object directory in kernel memory — two very advanced DKOM hiding techniques.



Figure 9-14:
Using
DarkSpy
to detect
BadRK
Demo's
driver.

Figure 9-15 illustrates using DarkSpy's Online Analyze Registry function to locate and remove the Service Registry autostart for BadRKDemo's driver (specified by `ControlSet1`). This is possible because DarkSpy's method of accessing the Registry doesn't rely on a system Registry call.



After you reboot, BadRKDemo's driver won't be loaded, so you can delete its file on disk. After you remove the driver file, BadRKDemo is completely removed (since it is the sole rootkit component).

In general, DarkSpy can detect the following techniques that most other anti-rootkit programs are incapable of detecting:

- ✓ DKOM of the Process List and PspCidTable (for example: FUTo)
- ✓ DKOM of driver list and object directory in kernel memory (for example: BadRKDemo)
- ✓ Kernel inline function hook of system Registry call (for example: BadRKDemo)
- ✓ Inline function hooks of kernel services that affect the file system (for example: VXK rootkit)

DarkSpy's authors have created a proprietary test version of DarkSpy that runs successfully on Windows Vista Beta 2 with all its current features supported (with the exception of hidden port detection). That bodes well.

GMER

GMER is a relatively new and very popular player in the anti-rootkit arena. Named after its Polish author, Przemyslaw Gmerek, GMER comes as a stand-alone executable file called `gmer.exe`. When run, `gmer.exe` self-extracts a device driver called `gmer.sys` and a DLL called `gmer.dll`. Both these files are required for GMER to function properly in kernel mode and to implement its system-monitoring feature (which we'll discuss in a minute); first, however, take a look at GMER's rootkit-detection features.

When you launch GMER, you'll only see a Rootkit tab and a tab marked with >>> in the main display. Clicking the >>> tab updates the display so you can view and access GMER's functions — which you can find on the following tabs:

- ✓ **Processes:** The Processes tab lists all the running processes in the system. If any hidden (“rooted”) processes are found, they're listed in red (as is the case in IceSword and DarkSpy).
- ✓ **Modules:** This tab lists all the kernel modules (device drivers) loaded into the memory.
- ✓ **Services:** The Services tab lists all Windows services present in the system; hidden services are shown in red.
- ✓ **Autostart:** The Autostart tab lists registry autostart locations for programs that start automatically when Windows loads.
- ✓ **Rootkit:** The Rootkit tab accesses one of the most important functions of GMER. Choosing the Scan option located in GMER's Rootkit functions allows you to scan your whole system for rootkit activity. At the conclusion of the scan, GMER produces a list of all the hidden resources it finds — such as files, folders, processes, threads, services, and Registry keys.
- ✓ **CMD:** The CMD tab provides an alternative environment to execute MS-DOS batch files and Registry scripts. This section is useful when the Registry Editor or the command prompt isn't functioning, or when the system is booted in GMER Safe Mode (more about that feature in a minute).
- ✓ **Settings:** The Settings tab provides various options that enable you to customize GMER's functions.
- ✓ **Log:** The Log tab works in conjunction with GMER's System Monitoring and Tracing — and that function must be activated for the Log function to work. If System Monitoring and Tracing is enabled, GMER logs all system events and displays them when you choose the Log tab.

Rootkit detection and removal using GMER

One chief advantage of GMER, as compared to many other anti-rootkit programs, is its automatic detection-and-warning feature. When GMER starts, it automatically checks key data structures that are known targets for rootkit modification. If it finds any changes that indicate hidden processes, drivers, or services may exist, GMER lists its findings and advises you to perform a full-system scan to detect all rooted components (we show you an example of this process a little later in this section).

To start a complete system scan, go to the Rootkit tab and specify the areas of the system you'd like to scan, and choose the Scan button. You can scan the following areas of your system:

- ✓ **System:** Scans the OS kernel components to check their integrity.
- ✓ **Devices:** Scans devices and their drivers to see whether any hidden rootkit drivers are masquerading as devices.
- ✓ **Processes:** Scans for hidden (rooted) processes and threads.
- ✓ **Libraries:** Scans for hidden Dynamic Link Libraries (DLLs).
- ✓ **Modules:** Scans for hidden drivers loaded in memory.
- ✓ **Services:** Scans the services registered in the system for rootkit services.
- ✓ **Registry:** Scans the Windows Registry for hidden rootkit entries.
- ✓ **Files:** Scans the file system, specifically for rooted files and folders.

After the scan is finished, any hidden items found during the scan are listed. You can right-click any item in the result listing and choose whether to delete or restore it, depending on the type of entity (or threat) it represents. For example, if a listed item is a hidden file, normally you'd delete it; but if the item is a *resource* that's hidden with an inline kernel hook, like BadRKdemo's driver was, then you should restore that hook before deleting the item. That's because you don't want the OS to continue to reference code in a rootkit driver that doesn't exist — doing so will result in a system crash.



Don't select the Show All option listed under the Rootkit tab before you perform a scan, or GMER will show all SSDT entries and drivers present in the system, instead of showing only hidden components and critical system modifications.

You can save the GMER rootkit-scan results by choosing Copy (to copy the results to the Clipboard), opening a Notepad file, and pasting the Clipboard contents into that file. You can post this saved log to any security forum listed in Chapter 13 if you need help interpreting the results.

Because GMER is both easy to operate and very effective, it has become a favorite rootkit-detection-and-removal tool. GMER was the first anti-rootkit program to both detect and remove the infamous pe386 rootkit. GMER's author, who's known as Gmer (it's his name abbreviated), makes every effort to update his program as soon as a new rootkit threat emerges.

Figure 9-16 illustrates how GMER immediately detects pe386 rootkit modification as soon as it starts (notice the WARNING!! window). GMER detects both the pe386 service and the location of its hidden module (driver) in the ADS stream of the system32 folder. It also detects pe386's SYSENTER hook. If GMER's ADS scanning option is enabled, GMER lists the first ten ADS streams it encounters (good or bad) during its scanning.




Figure 9-16:
GMER
Rootkit
Activity
Alert.

Figure 9-17 shows you how the pe386 service is removed: You highlight its service Registry entry and choose Delete the Service.




Figure 9-17:
Deleting a
rootkit
service.

GMER can detect processes and drivers (modules) hidden by FUTO, and it can also kill those same hidden processes. The GMER Web site has sample logs — even videos — of many prevalent rootkits that GMER detects and removes. (You can find those resources at www.gmer.net.) Here's a preview — a list of available logs and the rootkits they feature, as shown at the GMER Web site:

- ✓ **Gromozon Adware Rootkit:** www.gmer.net/gromozon.log
- ✓ **pe386 Rootkit (a.k.a. Agent Rustock):** www.gmer.net/rustock.log
- ✓ **Haxdoor Rootkit:** www.gmer.net/haxdoor.log
- ✓ **HackerDefender Rootkit:** www.gmer.net/hxdef.log
- ✓ **BadRKDemo:** www.gmer.net/badrkdemo.log

And here's where you can find a log that shows how GMER can detect a new rootkit called `phide_ex`, which was developed by the pe386 rootkit's author:

www.gmer.net/phide_ex.log

Enabling system monitoring and tracing in GMER

GMER can also be used as a Host Intrusion Prevention Software (HIPS). If the system-monitoring feature of GMER is enabled, it can intercept the actions such as creation of processes, services, and loading of driver, and prompt you to allow or deny the action. This feature is useful for blocking the installation of rootkits — or of any suspect or unknown programs until after you've researched them.



Sometimes PUPs (*potentially unwanted programs*) are bundled with beneficial programs. It never hurts to know what you're getting.

To enable system-monitoring, you need to check the options System Protection and Tracing and User's Application Protection and Tracing in GMER's Settings tab. Then you can select the different system areas and activities you'd like GMER to monitor. Here are your options:

- ✓ **Processes:** Monitors the creation of processes by user's applications.
- ✓ **Libraries:** Monitors the loading of DLLs by user's applications.
- ✓ **Drivers:** Monitors the loading of drivers by user's applications.
- ✓ **Files and folders:** Monitors the creation of executables by user applications.
- ✓ **Registry:** Monitors the Windows Registry for modification or creation of autostart entries by user applications.
- ✓ **Network:** Monitors the creation of unauthorized network connections attempted by Internet Explorer, Outlook, and Outlook Express.

Here's an example of how you can configure GMER to enable process monitoring on your system: If you select the Processes and Prompt Before Creating New Processes options in the Settings tab, GMER will intercept the creation of each process in the system — and request your permission before allowing a process to run. If GMER intercepts a new process creation, you're prompted to choose an action from the following list so GMER knows how to respond:

- ✓ **Allow:** Allows process creation.
- ✓ **Deny:** Blocks process creation.
- ✓ **Prompt:** Requests user approval every time a process is created.
- ✓ **Disable:** Disables or terminates the process.

If you know that the process being created is legitimate, then you can allow it to run. You can have GMER remember your decisions by clicking the Save button in the dialog box it displays when it prompts you for confirmation. GMER's learning mode is similar to the learning mode that firewalls use to create a list of programs that have been approved for network connections.

GMER Safe Mode

Sometimes it's desirable to have a *clean* system environment — in which only essential system processes are running — to facilitate the removal of rootkits and other tenacious infections. A clean environment makes it easier to track down and remove rootkits and other forms of malware (which often have measures in place to monitor and restore any infected items you remove).

By choosing the Safe button in the Processes tab, you instruct GMER to create this type of environment (which is much more restrictive than Windows Safe mode). Doing so causes GMER to prompt you for approval before initiating a reboot; when the computer restarts, only the `gmer.exe` and `csrss.exe` (system) processes will be running. The Windows Desktop and Windows Explorer (Windows GUI) will both be noticeably absent in GMER Safe Mode. Even so, you can still navigate through the file system and locate specific files or folders; GMER provides its own file-system browser. Choosing the Files button under the Processes tab opens a Window (similar to Windows Explorer) that enables you to view and delete files and folders. You can also choose the CMD tab to access the Registry-editing and command-console functions we describe next.

GMER Registry Editor and Command Prompt

If you need to add, edit, or delete any Registry keys or branches, then you can use GMER's built-in Registry feature. Choose the CMD tab and select the REGEDIT.EXE button. After that's done, you can write Registry commands or scripts in the upper black command box and click the Run button to execute them.

Similarly, by selecting the CMD.EXE button within the same CMD function, you can execute MS-DOS commands or batch scripts. The output of any commands, batches, or scripts (Registry or DOS) you execute is displayed in the lower black window labeled `Log`. If you decide you'd like to enlist another program to assist you in your troubleshooting efforts, here's how:

- 1. Select the Processes tab and choose the Command option.**
- 2. Type the path and program executable (EXE) of the program you'd like to start and then click the Run button.**
GMER will start the program for you.
- 3. When your ready to restart the system, click the Restart button in the Processes tab.**

These three options (Files, Regedit, and CMD) are very useful when you're working in GMER Safe Mode — or if you have a severely infected system where the Registry Editor, Command Prompt, Windows Explorer, and Task Manager are no longer functional because they've been disabled by malware.



We're sure you'll be happy to know that GMER anti-rootkit is included on your DART CD.

Detection for 64-bit Windows platforms

We've had inquiries regarding the existence of rootkit tools that can be run on 64-bit Windows 2000, XP, and Server 2003 Platforms. Very recently, Resplendence updated their RootKit Hook Analyzer to provide support for all these platforms. The RootKit Hook Analyzer lists SSDT hooks and resolves them back to the program that created them. The program is free, and runs on 32- and 64-bit editions of the platforms we've listed here. Resplendence RootKit Hook Analyzer can be downloaded at this link:

www.resplendence.com/hookanalyzer

Detecting Keyloggers

It's a common practice for rootkits to install backdoors to transmit sensitive information back to a remote attacker. This information can consist of passwords, account information, client lists, or other confidential data with a common goal in mind — exploiting the information gathered for financial gain.

In Chapter 8, we covered how hackers can install sniffers to harvest valuable confidential information that's transmitted over a network. Now we'd like to address a variation of this threat that's more commonly found on standalone computers: *keyloggers* — programs that secretly record everything a user types.

Keyloggers can be installed for a variety of reasons — some illegal, some for the purpose of monitoring employee or child behavior. But whatever the reason, the person being spied on resents (or at least never appreciates) the intrusion — especially if the ultimate purpose of the keylogger is identity theft or some other form of financial exploitation.

Types of keyloggers

There are two types of software keyloggers: Those that operate at the user level using a DLL injection, and those that achieve their functionality by installing a device driver. Regardless of their mode of operation, keyloggers work the same way — by capturing every user keypress and then recording it to a file. *Backdoor keyloggers* then transmit the collected information remotely to blackhats who are eagerly waiting to get their hands on something juicy. DLL-based keyloggers hook Windows messages to intercept and

capture keyboard input. Windows messaging is used to facilitate communication between the OS, devices, and processes running on a computer. As is the often case — and it applies to Windows Messaging — the same beneficial feature that fulfills a useful purpose can be exploited by blackhats with malicious intent. DLL-based keyloggers can even capture passwords that were filled in using an application's auto-complete feature, but this requires hooking and filtering the SetWindowsHookEx API, as opposed to using a message hook.



There are many different Windows messages that exist for various purposes, but the one used specifically to transmit keyboard information is called `WM_KEYBOARD`.

Detecting keyloggers with IceSword

We can use the Message Hooks function that IceSword provides to easily detect DLL-based keyloggers. Choosing the Message Hooks function causes IceSword to list all message hooks present in the system along with the process responsible for hooking the message. Some of the hooks listed may belong to legitimate applications such as your firewall, antivirus program, and even IceSword. However, because many keylogger programs use the `WH_KEYBOARD` message hook to intercept `WM_KEYBOARD` messages, you can look for that particular hook; it's a dead giveaway.


To monitor your system for DLL-based keyloggers, you can download and install IceSword (see Bonus Chapter 1 for more information about that program, including its download link). Once you've installed IceSword, follow these steps:

- 1. Open IceSword and choose the Message Hooks toolbar icon.**
- 2. Examine all message hooks and see whether any listed under the Type column use the `WH_KEYBOARD` hook.**
- 3. If you find any entries using the `WH_KEYBOARD` hook, then look under the Process Path column of the same entry.**

That's where you can find the folder location of the keylogger's process.

Figure 9-18 shows IceSword detecting a DLL-based keylogger called `keylogger.exe` that uses `WH_KEYBOARD` hooking. Notice the Process Path field identifies the location of the keylogger executable file on disk. This particular keylogger is a freeware product, not known to be associated with a rootkit or a backdoor (we've used it for illustration purposes) — but there's a clear advantage to using IceSword to detect keyloggers that intercept Windows messages: Even if the keylogger was hidden by a rootkit, it would still be made visible by using IceSword's Message Hooks function.


Figure 9-18:
Using
IceSword's
Message
Hooks
function to
detect
keyloggers.



Detecting keyloggers with Process Explorer

Figure 9-19 illustrates Process Explorer's ability to trace the keylogger's DLL that has been injected into the `Explorer.exe` process. Notice that when we select `Explorer.exe` in Process Explorer's upper pane, the DLLs loaded by that process are listed in Process Explorer's lower pane: The `keylogger.dll` is clearly identified; though it has no company name, description, or version number, its filename is a dead giveaway.

Figure 9-19:
Locating a
keylogger-
injected
DLL with
Process
Explorer.



There is a reason that `Explorer.exe` was the targeted process that `KeyLogger.DLL` injected itself into. It was specifically chosen because it's the parent process of all processes launched by the user. Therefore any user programs running below that process are also injected with `KeyLogger.DLL`. Even if this keylogger were cloaked by a rootkit, you could find it by running Process Explorer in combination with `AntiHookExec`, as described earlier in this chapter.

Also as we mentioned earlier, keyloggers can be DLL-based or driver-based. If a keylogger is driver-based, you can use Process Explorer to identify the keylogger driver. Here's how that works:

- 1. Open Process Explorer and choose Options ⇨ Verify Image Signatures.**
- 2. In the upper pane's Process Tree display, double-click the System process.**

The lower pane display gradually updates to reveal all drivers present on your system — sorted by the Company Name field.

Because we chose the Verify Image Signatures option, all listed drivers have undergone digital-signature verification (as you can see by inspecting the Company Name field) — and the unverified drivers are listed first. The commercial keylogger driver may be here with no verification (and no company name), but most likely it will be identifiable by its name if it's a commercial product — as is the case here.

- 3. When you find the keylogger entry, right-click it and choose Properties.**

This is how you identify the image path (location of the driver file) that takes you right to the keylogger driver's lair.

Figure 9-20 shows Process Explorer detecting a commercial keylogger driver called `iks.sys` (again, this keylogger is not known to be rootkit or backdoor associated, and was used for illustration purposes only). This particular keylogger program allows you to rename the driver file (pretty stealthy) — so the lesson is: don't expect a keylogger to follow a consistent naming convention that aids in its identification and removal.

An alternative method of detecting driver-based keyloggers involves checking the drivers associated with your keyboard device, using the Keyboard function in the Windows Control Panel. To do so, open the Control Panel and select Keyboard ⇨ Hardware ⇨ Properties ⇨ Driver ⇨ Driver Details. You'll see a Driver Files Details window that lists all keyboard drivers on your system. Signed Microsoft drivers show green check marks to the left of them. If any drivers listed are missing that check mark, then you may have a kernel-mode keylogger on your hands.




Figure 9-20:
Using
Process
Explorer to
detect a
keylogger
driver.

Because keyloggers (as mentioned earlier) are often in cahoots with backdoors, let's go full circle and examine how a backdoor trojan — a.k.a. RAT (remote-administration trojan) — can transmit keyboard information (and so much more) from a compromised system. We expose the underbelly of that unscrupulous practice in our next section.

Tracking a RAT: Using Port Explorer to trace Netbus 1.60

The most malicious blackhat rootkits install covert backdoors — typically to steal and transfer information over extended periods of time. Often the fate of your rootkit-infected system is decided by the extent of the damage that the rootkit-accompanied threat has inflicted. That damage is of far greater importance than the practical matter of whether the rootkit itself can be removed safely.

The most insidious backdoors are remote-administration trojans (RATs) under remote human control. A RAT effectively reduces your system to nothing more than a puppet, subject to the whims of a remote hacker. In this section, we examine how a port-to-process mapping program can play a critical role in identifying and disinfecting a well-known RAT called Netbus.

Although Netbus has no rootkit characteristics per se, the method we apply here can be used to track down a rootkit-related RAT after the rootkit that was hiding it is unclocked. The port-to-process mapping program we use in the example is Port Explorer, which you can run from the command line in combination with AntiHookExec. That combination will reveal ports hidden by user-mode rootkits — and the processes that own them.

In this section, we show you how you can use Port Explorer to track down Netbus 1.60 — one of the best-known and most-studied trojans. Netbus was one of the very first remote-administration trojans to infect Windows platforms. Its commands are issued in plain text language so it's an ideal candidate for us to study. Netbus installs a server (slave) program on its infected target and a client (master) program on the computer that controls it. These two program components are required for the hacker to communicate with the trojan.

Figure 9-21 shows the main display of Port Explorer. As shown in the figure, the Netbus trojan process is `c:\windows\patch.exe`, and it has two ports open: 12345 and 12346. Its process ID is 3984.

Figure 9-21:
Port Explorer showing a trojan process — `patch.exe`.

Process	ProcessID	Protocol	Local Address	Local Port	Remote Address	Remote Port	Status	Sent	Receive
* SYSTEM	4	TCP	0.0.0.0	445	0.0.0.0	0	LISTENING	---	---
* SYSTEM	4	TCP	10.1.1.2	139	0.0.0.0	0	LISTENING	---	---
* SYSTEM	4	TCP	10.1.1.2	137	***.*	*	LISTENING	---	---
* SYSTEM	4	UDP	10.1.1.2	138	***.*	*	LISTENING	---	---
* SYSTEM	4	UDP	0.0.0.0	445	***.*	*	LISTENING	---	---
lsass.exe	656	UDP	0.0.0.0	500	***.*	*	LISTENING	0.00K/s (0.00K Total)	0.00K/s
lsass.exe	656	UDP	0.0.0.0	4500	***.*	*	LISTENING	0.00K/s (0.00K Total)	0.00K/s
svchost.exe	884	TCP	0.0.0.0	135	0.0.0.0	0	LISTENING	0.00K/s (0.00K Total)	0.00K/s
svchost.exe	1068	UDP	0.0.0.0	1039	10.1.1.1	53	LISTENING	0.00K/s (0.08K Total)	0.00K/s
svchost.exe	1068	UDP	0.0.0.0	1040	10.1.1.1	53	LISTENING	0.00K/s (0.08K Total)	0.00K/s
svchost.exe	1088	UDP	127.0.0.1	1900	***.*	*	LISTENING	0.00K/s (0.00K Total)	0.00K/s
svchost.exe	1088	UDP	10.1.1.2	1900	***.*	*	LISTENING	0.00K/s (0.00K Total)	0.00K/s
svchost.exe	1088	TCP	0.0.0.0	2869	0.0.0.0	0	LISTENING	---	---
patch.exe	1168	TCP	0.0.0.0	12345	0.0.0.0	0	LISTENING	0.00K/s (0.00K Total)	0.00K/s
patch.exe	1168	TCP	0.0.0.0	12346	0.0.0.0	0	LISTENING	0.00K/s (0.00K Total)	0.00K/s

#	Time	PID:Process	Action	Protocol	Local Address	Remote Address	Status	Bytes	Country
14	12:57:42	1168:patch.exe	LISTEN	TCP	0.0.0.0:12346	0.0.0.0:0	SUCCESS	0	
13	12:57:42	1168:patch.exe	OPEN	TCP	0.0.0.0:0	0.0.0.0:0	SUCCESS	0	
12	12:57:42	1168:patch.exe	LISTEN	TCP	0.0.0.0:12345	0.0.0.0:0	SUCCESS	0	
11	12:57:42	1168:patch.exe	OPEN	TCP	0.0.0.0:0	0.0.0.0:0	SUCCESS	0	


15 Sockets (13 System, 2 Hidden, 0 Normal) Admin Privilege Enabled WinXP

Netbus's primary port is 12345, which is its *control port* — that's the port where the trojan receives the master program's incoming commands as it lies in wait on the infected host. If the master program sends a command that involves sending a file (say, sending a screenshot of the user's desktop), the trojan slave program responds by sending the requested information on Port 12346 — in effect, Netbus's file-transfer port.

Port Explorer uses a color-coded system to simplify its display. It displays Netbus's sockets in red because Netbus is running as a hidden process (to Port Explorer hidden means it has no visible desktop components, even though its process is not necessarily hidden in the Active Process list). Virtually all RATs show up in red for this reason — because they're not prone to announcing themselves on the desktop. Port Explorer uses black text to display processes that do have visual desktop components, and system services show up as blue.

Here's a bit of revenge: Port Explorer can be used to reverse-engineer the trojan's own communications protocol — to the point where remote disinfection is possible. To do that, you must enable spying on the `patch.exe` process. Then any data packets transmitted to (or by) `patch.exe` will be captured for analysis. You enable spying by choosing Utilities ⇄ Socket Spy on Port Explorer's menu. As is shown in Figure 9-22, when the Socket Spy Packet Sniffer window opens, you add the `patch.exe` process to the Socket Spy list of processes to spy on by entering `patch.exe`'s PID into the open box and choosing Add PID. From that point on, all network communications to and from `patch.exe` are logged by Port Explorer — and all data packets transmitted will be captured for analysis.

Figure 9-22:
Enabling
spying on a
trojan
process.



Every trojan has its own “language” (or protocol) used by the master program and the slave (the trojan). The Netbus trojan essentially identifies itself to the program that has connected to it, regardless of what that program is. Normally, the connecting program would be the “master” program that connects to the trojan from a remote location to control it.

Figure 9-23 shows a captured packet of data. It's the first packet that was captured — as indicated by its `Index # 0` Index number. Its destination is Out, which means the `patch.exe` sent the data, as opposed to receiving it. You can also see that the data it sent was `Netbus 1.60 -`. Aha! This is the command that identifies Netbus to any program that connects to it.




Figure 9-23:
Examining a
packet sent
by the
trojan.

If any program connects to the Netbus trojan on TCP port 12345, the trojan immediately reacts by sending out the `Netbus 1.60-` string. Normally the connecting program would be the master program, and this command would notify the master program that it has indeed established a connection to the trojan. Some trojans remain silent until the master sends a please-identify-yourself type of command, but Netbus just speaks right up; it's a good example of a trojan that immediately identifies itself to the program that has connected to it — making remote detection very easy.

Many backdoor programs have elaborate master control programs with all the amenities — even a nicely designed graphical user interface (GUI). This makes them simple and easy to operate. (Blackhats like an easy ride, too.)

Figure 9-24 shows the Netbus master control program that the hacker uses to remotely connect and send commands to the trojan to control it. This image illustrates what the hacker sees on the client side.




Figure 9-24:
The Netbus
master
control
program.

Notice all the functionality the trojan's master control program has available to it. It allows its controller to open the CD-ROM, show an image, get a screen dump, control the mouse and sound, and even go to a URL. It's really pretty alarming to see how much control a remote hacker can establish over your computer or server.

Clicking Server Admin brings up the server-control functions. Notice that the Remove Server button is highlighted. Clicking Remove Server, as we did here, disinfects the trojan program on the infected client computer.

By using Port Explorer's Socket Spy utility we can easily see exactly what the master program sent to the slave — in this case, a packet that disinfects it, as shown in Figure 9-25, which reveals the data sent in the captured packet. The command `RemoveServer;1`, followed by a carriage-return character (`0x0D`), performs the *coup de grâce*. Armed with that knowledge, a programmer could easily write a remote-disinfection program — one that simply sends that command to any system infected with that trojan.




Figure 9-25:
Examining
the Netbus
disinfection
command.

Because the master control panel that sends the disinfection command is under the control of the attacker on the remote client computer, you may be wondering how you can accomplish disinfection from the infected computer (server). The answer is: Be sneaky. Send out the disinfection command via Netstat port 12345 — either programmatically or via Telnet.




Telnet is a Windows program that enables computers on a local area network (LAN) or the Internet to communicate. It's a console application that runs from the Windows command prompt. The Telnet program executable (`telnet.exe`), resides in the Windows system32 directory.

The next order of business is to illustrate two handy methods of disinfection. The first uses Windows Telnet capability (from the command prompt) to connect and issue the disinfect command to the master program. This is all done from the infected computer. The second method uses a very simple Visual Basic script to send out the disinfect command. Read and enjoy.

Method 1: Using Telnet to send the disinfection command

First, you open a command prompt window by clicking Start and choosing Run, typing **cmd** into the Open: field, and then clicking OK or pressing Enter. Figure 9-26 illustrates the command console and the commands that are sent via Telnet. As you can see, disinfection is quite a simple process.

Figure 9-26:
Sending the
disinfect
command
using Telnet.



```
C:\WINDOWS\system32\cmd.exe
C:\WINDOWS>telnet 127.0.0.1 12345
Connecting to 127.0.0.1...
NetBus 1.60
RemoveServer;1
Connection to host lost.
C:\WINDOWS>telnet 127.0.0.1 12345
Connecting to 127.0.0.1...Could not open connection to the host, on port 12345:
Connect failed
C:\WINDOWS>
```

The first line in Figure 9-26 shows

```
telnet 127.0.0.1 12345
```

This is the command we use to start the Telnet program and to pass it two parameters: **127.0.0.1** and **12345**. The first parameter (127.0.0.1) represents the computer to connect to; the second parameter (12345) represents the *port* we want to connect to — the Netstat command port.

The second line shows `Netbus 1.60`. We didn't type or send this command — that's what the trojan server sent back to acknowledge it had received the first command.

Line 3 shows `RemoveServer;1` — and this represents the disinfection command. Because the trojan acknowledged it's ready, we can now enter this command by typing it exactly as shown on the third line, and pressing Enter to send it.

Line 4 displays the message `Connection to host lost`. This message was sent by the Telnet program to inform us that the connection had been terminated. Because we didn't terminate it, we know that the trojan server initiated the termination of the connection, as anticipated.

Line 5 is the same command we sent in Line 1, which attempts to establish a connection with the trojan. We tried to reconnect, but this time it failed because the trojan has now removed itself, proving that we've been successful at remotely disinfecting the Netbus trojan.



It's possible that the remote attacker may have configured Netbus to use a password — in that case, an additional command would have to be inserted into the sequence given here, to bypass that password. The command that takes care of that minor imposition is `Password;1;` — and it would have to be executed before the `RemoveServer;1` command in Line 3.

Method 2: Writing a simple Visual Basic program to disinfect Netbus:

The following is a small program coded in Visual Basic that accomplishes the same disinfection task described in Method 1:

```
Private Sub Main ()
Winsock1.Connect "127.0.0.1", "12345"
End Sub

Private Sub Winsock1_Connect ()
Winsock1.SendData "RemoveServer;1" & Chr(&h0D)
Winsock1.Close
Msgbox "Successfully sent the Netbus disinfection
      command."
End Sub
```

Note: In the Visual Basic code given here, `Winsock1` is the Microsoft Winsock Control. The Microsoft Winsock Control must be added to the Visual Basic project (choose Projects from the Components menu in the Visual Basic IDE), and its name by default is `Winsock1`. (If you add a second one, it will be named `Winsock2`, and so on.) This allows the Winsock control to be used as an object; you can (for example) insert commands such as `Winsock1.Connect 127.0.0.1 12345` because `Connect` is one of the functions that Winsock Control offers.

The beauty of the remote-disinfection technique is that it can be used to disinfect any number of remote computers connected to a network. You don't have to visit every infected workstation to eradicate the threat — you can simply use a small program that sends the disinfect command to all infected remote computers to eliminate Netbus quickly and efficiently.



Even though Netbus is not installed with a rootkit, many rootkits do hide RATs that behave in a manner very similar to that of Netbus. You can run Port Explorer in combination with AntiHookExec to reveal any user-level rootkit hidden trojan processes that maintain open ports. Assuming you have AntiHookExec installed on your system, you can simply invoke Port Explorer from the run line as follows:

1. Click Start and choose Run.
2. In the Open: field, type in the following line:

```
AntiHookExec "C:\Program Files\Port Explorer\  
PortExplorer.exe"
```

3. Press Enter or Click OK.

Part IV

Readying for Recovery

The 5th Wave

By Rich Tennant



In this part . . .

Now is the time to step out of the theoretical and into the practical. So here's where we put you in the driver's seat, while describing the whole process of dealing with a rootkit intrusion. In these chapters, we emphasize rootkit readiness — but we also help you decide what to do if the worst-case scenario becomes reality. We guide your preparations for a competent recovery in case your computer or network is breached. We also give you pointers on safely preserving essential forensic evidence should you decide to pursue legal channels against your intruder. We discuss the best way to treat a compromised system to both safeguard evidence and minimize damage to your important files.

We show how to clean up after an intrusion — review the different types of rootkit threats you may be facing, and spotlight the best methods for addressing them. Virtual soldiers need to know how and when to retreat, so we conclude with a look at how to know when it's time to wipe that hard drive and start from scratch — and provide guidelines for getting that done, should it be necessary.

Chapter 10

Infected! Coping with Collateral Damage

In This Chapter...

- ▶ Choosing a forensic path
 - ▶ Pulling the plug
 - ▶ Deciding the computer's fate
-

This chapter is all about damage control — what you need to do after a rootkit compromise, both on single computers and on network servers. We have already shown techniques for detection and removal of rootkits. Rootkits may be used to promote criminal activities that can be traced back to your computer. The villain gets away and you're left holding the bag. Authorities may end up investigating you — and perhaps even assume that you're a blackhat hacker. (Can you say, "Adding insult to injury?")

Deciding What to Do if You're Infected

The question of what to do after your system has been hacked depends on whether you're an individual user or if you're part of a network. For rootkits, backdoors and remote-access trojans (RATs), the results are moot. No way yet exists to be absolutely sure you've removed all of them. Mitigating collateral damage is not going to be very important if you need to reformat and reinstall. If you're on a network, you need to report the condition of your workstation to your Network Administrator immediately. The importance of doing so will assist in checking the other workstations in the network for possible compromise. If law enforcement needs to be contacted, it's important to isolate all affected systems in order to gather evidence.



Some computing reports advise that staying online when you know you've been hacked helps gather more evidence on a blackhat hacker — but this can lead to further loss of physical evidence. The less the affected computers are used, the more evidence is available to be collected. That's because data files are notoriously variable and changeable. When a computer is compromised, it's a good idea to make a copy of the entire contents of its hard drive, on separate media, so you can compare the copy with your clean and normal backups. Isolate the computer or workstation and disconnect it from the Internet first.

When making copies of hard drives for possible use in court, you must use special imaging tools that do not alter the state of the files in any way. Here's a handy pair of sites where you can obtain free or commercial versions of those tools:

```
www.forensics.nl/toolkits
```

```
www.opensourceforensics.org/tools/windows.html
```

For commercial tools, EnCase is the standard. EnCase is discussed in the section “Getting a pro to analyze the evidence” later in this chapter.



Imaging and file-backup programs are included on the *Rootkits For Dummies* DART CD. Imaging programs help you to make full backups of your system and data files for your entire hard drive(s). File-backup programs such as The Replicator can be set to continuously maintain backups of important and sensitive files.

When you're dealing with a rootkit compromise, it's especially important to have your clean backups on separate media, because your system can no longer be trusted. Backups can also help by providing a clean and normal state for comparison. After removing the rootkit and its accompanying malware, you may still have to reformat the hard drive and reinstall your operating system before you restore your backups.

If your scanner finds malware or viruses, don't panic. If the malware is in a non-system file that you can live without, then delete it. If it's in a system or cherished program folder, quarantine it. Antivirus, anti-trojan, and anti-spyware applications have folders for quarantining detected files, and usually that's where the pests go automatically. If malware damages your system files — Windows File Protection (as discussed in Chapter 7) automatically kicks in — and replaces them with authentic new copies. You can then leave the corrupted file in quarantine (for later careful examination) or delete it — it will no longer be infectious.

Many anti-malware product vendors do provide file-analysis services that you can take advantage of if you're unsure whether a particular file is infected. You can contact them and follow their instructions. Better yet — if you want immediate feedback — just upload the infected file(s) to the Virus Total or Jotti online scanners (which are discussed in Chapter 9) and you'll receive a definitive report within minutes. It's a good idea is to keep backup copies of all your cherished programs on separate media (such as CD/DVDs) in case any of them become corrupted by malware.



If you're stumped about what to do with your malware detections — or if you suspect you have a rootkit — please visit one of the reputable security-forum Web sites listed in Chapter 13. You can ask for help from real experts.

Normally the performance of your computer or network gives you clues to the possible presence of a rootkit or an intruder. On a network, you get performance indicators such as excessive usage of bandwidth, CPU cycles, memory capacity, or disk space. All the methods discussed in Chapter 8, “Sniffing Out Rootkits,” should be employed to compare current network activity to normal activity. Similar clues will appear on single-user computers, due to increased usage of system resources. Whether you're on a network or single-user system, you may also experience system crashes and instability. One of the most prevalent indicators of a rootkit infection is the persistence of obvious operational and performance issues that continually linger despite clean traditional scan reports.

Knowing when to give up and start from scratch

It's often thought that once a rootkit's detected on your system that it's the only “cure” is to reformat your hard drive and reinstall your operating system. We're skeptical of any hard-and-fast general rule that supposedly applies to all situations. The final verdict definitely depends on how dangerous the threat (or threats) that the rootkit was hiding — and your computer was harboring — are found to be.

If the rootkit was hiding a payload that was more nuisance than threat (such as your garden-variety adware or spyware program), there's no need to reformat and reinstall. For example, the Apropos rootkit could be removed without any far-reaching repercussions, and also the Vundo trojan (responsible for WinAntispyware and Antivirus Pro pop-ups) can be installed with a rootkit, but once the rootkit is unclocked, removing the trojan is no different from any other Vundo trojan removal. The consequences are ultimately identical — and no permanent damage was done.



Sometimes it may be necessary to reformat and reinstall your operating system merely because of an attempted rootkit removal that didn't go as planned. That's because kernel-mode rootkits are intimately bound to kernel functionality; suddenly removing a kernel driver could cause your computer to become unstable and unbootable. The best general approach to follow is to first remove the rootkit service's autostart in the Registry, and then reboot so the rootkit driver is no longer loaded — and then proceed with removal.

Usually you can research the particular threat you have on your system to determine if removal can be accomplished gracefully. Unlike trojans in general, new rootkits are not introduced that frequently. Often you can find adequate removal instructions if you perform a Google search on the files and processes that the rootkit detectors uncover. You can always seek assistance from one of the online security forums; getting advice from someone who may be more experienced at rootkit detection and removal can't hurt.

Yet there are no hard or fast rules. If you're on a standalone computer, then you can be somewhat more flexible in your decision-making because the infection is confined to a single PC. You can try removing the rootkit and its associated malware payload, and then reassess your system to see whether the rootkit symptoms have subsided. If you're on a network, and are proficient at malware removal, you may decide to attempt removing the rootkit and any infected files from the server (and any affected networked computers). If the rootkit is memory-based (non-persistent) — which means it can't survive a reboot — then rebooting and disinfecting the network may suffice. Even kernel-mode rootkits can be safely removed if they're hiding a known — and *contained* — type of threat for which proven removal methods exist. On the other hand, if your network or computer is the victim of a targeted backdoor attack — and the full scope of the damage is extensive and difficult to assess — then reinstalling the operating system is probably the best course of action.

Certainly the most devastating rootkits are those that shield a remote-administration trojan (RAT) and its ongoing activities. If your computer or network has a RAT installed that has effectively opened a backdoor for some indeterminate length of time, then your computer may be completely compromised. This situation usually requires that you start from scratch, depending on whether the RAT had set up shop on your computer or continually sent back sensitive information or passwords to a remote hacker. For security reasons, that situation would most likely merit starting from scratch, unless you could accurately assess the extent of all damage on the infected system(s). Of course, even then you would have to change all passwords and notify any institutions or parties who were affected by the compromise. If you had adequate controls in place that enabled you to gauge the extent of the damage — say, offline databases and other system-baseline logs — then you could possibly determine how far the threat had penetrated and

accurately evaluate the damage it had inflicted. The bottom line is *safety and security*; unless you are absolutely sure that you can completely clean up a RAT — and the devastation it may have wreaked (which is unlikely) — then reformat-and-reinstall is the name of the game.

What happens when the patient can't be saved

If the prospect of completing rootkit removal seems too daunting or not worth the effort expended, then you may want to consider reformatting and reinstalling your operating system. In fact, many network admins who are confronted with the prospect of rootkit removal opt for the drastic approach: back up their data, sever all connections, reformat the hard drive, and reinstall the operating system. If a kernel rootkit (which can modify operating system data structures and possibly even kernel objects in memory) is present, then this approach might be the only way to kill it off. That's because removing a kernel rootkit is a delicate operation that can render the operating system nonfunctional, if it is done incorrectly. Furthermore, it's possible that the rootkit's far-reaching tentacles may have corrupted other OS system files as well (like their original UNIX predecessors). Of course, the entire recovery process is less painful if you happen to have a clean backup image of the hard drive on hand — and can use it to restore most of what you may have lost. You're best off if you've developed effective security habits such as systematically backing up data in a routine and timely manner — which would also greatly simplify a recovery operation.

The bottom line is, if you can no longer trust the integrity of your operating system, or even the data it contains, then starting from scratch may be the only answer.



This is why implementation of prevention measures is so critical. A full backup-and-recovery plan should be set in place in the event that a worst-case scenario does unfold. Hope for the best, but plan for the worst (which Chapter 11 can help you do).

Do you want to track down the rootkit-er, or just recover?

If you discover your computer or server has been infected by a rootkit or that you've suffered a break-in, one of the first things to consider is whether you want to track down the rootkit perpetrator or just recover — because

this decision has a great impact on what can be done on your system. If you know you've experienced a financial loss (or potential financial loss) due to the theft of confidential information, you may very well decide to collect evidence for litigation — and if that's the case, then you don't want to touch the computer after you realize it's infected.

No, you can't catch anything by touching a rootkit-infected computer — but you might ruin some evidence that could lead you to the rootkit-er. Of course, if you decide not to track down the rootkit-er, then fiddling with the evidence doesn't matter, and you can focus on trying to root out the rootkit or (as Chapter 11 describes) get started on a reinstall.

If the computer in question is a standalone machine, you may want to collect evidence and forward it to the appropriate agencies; it may be too much of a financial and emotional drain to pursue litigation. Often the criminals responsible for cyber-crime live outside the jurisdiction of the U.S. courts, and even if you're able to track them down, pursuing prosecution may be next to impossible.

Taking measured action

If your computer has *not* been hacked or broken into, but you still feel victimized by nuisance software programs that you feel were installed on your computer without your permission, you may want to pursue some less drastic avenues.

Most often rootkits have been used for illegal spying and monetary gain, but some are still being used by corporations seeking total control over what the consumer can do with their products (music or videos, for example). Complain to the company and/or vote with your wallet by refusing to purchase their goods. Corporations that have used these techniques are fast discovering that once consumers realize how they're being treated, sales plummet. The smarter companies will get the message. For example, you can believe that Sony got a flood of mail, phone calls, and e-mail from unhappy customers after their rootkit snafu. If a legitimate company wants to stay in business, enough of its customers complaining can — and in Sony's case, did — have an effect on company policy: they abandoned the distribution of DRM software (including the rootkit) on their music CDs and also offered their affected customers a choice of \$7.50 and one free music album download — or three free music album downloads. There are cases in the United States where adware or spyware companies have been sued for installing unwanted software on peoples' computers. Sometimes, when enough people complain to a legislative body or government official, an investigation can ensue. There are security researchers and advocates who will often take up the cause. They may supply compelling evidence obtained from conducting their own research, which can have a huge impact on the outcome of any case. The accompanying sidebar illustrates what's possible.

New York says, “Hands off my hard drive!”

The State of New York is quite aggressive in their intolerance toward incidents of unwanted installations. Here’s the gist of two recent lawsuits that were prompted by a multitude of user complaints — and supported with evidence supplied by security researchers.

In April 2005, New York Attorney General Elliott Spitzer (who has since become New York’s Governor — You Go, Elliott!!) sued Intermix Media (formerly eUniverse) for bundling spyware programs such as KeenValue and IncrediFind with its free games and screensavers. The suit alleged that this was done without sufficient notification — and it also alleged the bundled spyware programs were difficult to uninstall.

Also in April 2005, Elliot Spitzer brought suit against the adware company Direct Revenue LLC, for bundling pop-up advertising programs with their online offerings, and secretly installing unwanted software onto users’ computers by using drive-by-download techniques. Some of the programs that Direct Revenue distributed were VX2 BetterInternet (a misnomer if ever there was one!), Aurora, `nail.exe`, and OfferOptimizer.

To read more about the specifics of this case, you can refer to *People of the State of New York v. Direct Revenue, LLC — Documents and Analysis* by security researcher and Harvard University Doctoral candidate and law school student, Ben Edelman. Here’s where to find it:

www.benedelman.org/spyware/nyag-dr/

To be perfectly clear, neither InterMix nor Direct Revenue installed rootkits with their adware programs, but that doesn’t mean adware companies are not doing so. The adware company ContextPlus, Inc. distributed the infamous Apropos rootkit with its adware program, even though it eventually elected to suspend distribution of the whole shebang, reportedly by its own volition. There are numerous other examples of rootkit adware — CommonName, Elite Toolbar, and the very prevalent Gromozon rootkit (bundled with LinkOptimizer adware), to mention a few more.

If you find that your hard drive has an unwelcome guest, Malware Complaints is an online forum that deals with user complaints concerning spyware and adware programs.

www.malwarecomplaints.info/index.php

According to the Malware Complaint Web site, after collecting and organizing the complaints, Malware Complaints might

- ✓ Issue petitions, which can be offered to governments and official bureaus that deal with Internet security, thus making it clear to those officials how extensive the problems with malware have evolved.
- ✓ Make malware issues known to the news media — which, in turn, can make the malware and its makers known to the world — in an uncomplimentary light.

Another avenue you can pursue is to report your grievances to government agencies such as your state's Attorney General's Office or the Federal Trade Commission's Bureau of Consumer Protection if you live in the United States.

“My Computer Did What?!”

If you're quite certain your computer or network has been compromised resulting in a possible loss of financial or sensitive confidential information, or you suspect your computer has been used for illicit activity, it's important to take some data recovery steps before it's too late. This way you may be able to catch the intruders and limit your liabilities for any damages done. If your computer has been used to promote criminal actions ignorance may not be your best defense. It's possible that you could be held responsible for what was done with your computer by a blackhat hacker, whether you were aware of it or not. That's why knowing how to collect evidence in the proper manner is so critical.

These three rules are wise to follow when dealing with a hacked computer or network:

- ✔ Always treat your investigation as if it's going to end up in litigation.
- ✔ Always make “forensically sound” backups of all systems — and work only from those backups (no, we're not talking about CSI here, just something like it, and we'll fill you in on the exact details in the next section).
- ✔ If you expect legal proceedings, you should consult a computer forensics expert on how to best preserve evidence for use in court.

In this section, we provide advice and tips to help you deal with incidents where criminal activity may have occurred. We outline how to preserve and record evidence that may be useful to an investigation, both for you and for possible litigations. Having an accurate record of the incident can certainly assist in any inquiry.

Saving evidence to reduce your liability

Whether you decide to file a criminal suit or not, you should still try to save evidence to reduce your liability. In the event your identity was stolen, your bank account emptied, or your confidentiality violated — even if you have no one to prosecute or don't know the identity of your perpetrator — it's always wise to retain proof that you were victimized (you may need it later).

In either case, creating an exact image of the hacked system's hard drive makes sense. This way you have the evidence to fall back on, should the need arise.



You'll effectively lose anything in your computer's random-access memory (RAM) if you shut off your computer or reboot, so it's important to copy the data contained in RAM to an alternate medium before shutting down. This should be done on an active suspect system *before* copying the hard drive image.

Because it's not uncommon for a computer to have 1 GB or more of RAM these days, a CD that holds only about 750 MB is not a suitable medium for dumping a 1 GB RAM image to disc. The ideal medium for this kind of copying is a jump drive or portable USB Flash drive. Here's why:



✔ **It's big enough.** Flash Drives vary in their storage capacity, but today a 2GB flash drive can be purchased for around \$50. The data-collection process also requires an external hard drive having a capacity that exceeds the hard drive of the suspect computer system. A 300GB USB external drive can be purchased for \$150 – \$200, so having one available (should you need to use it) may be well worth the investment.

Always purchase a USB Flash drive with a capacity that exceeds the capacity of your RAM. A 2GB USB Flash drive has room to spare for most computers.

✔ **No driver hassle.** A portable USB Flash drive is immediately recognized by your computer when plugged in; it requires no software-driver installation.

✔ **It's easy to find when plugged in.** When you plug a USB Flash drive into an available USB port it's immediately assigned a letter, which indicates its device name. For example, if your CD-ROM drive is device D:, then your USB drive may be assigned device E.

Here's an outline that summarizes the steps you would take to preserve the state of your computer for later analysis:

1. **Obtain evidence from any currently open programs by taking photos or videos of the screen and the computer's setup.**
2. **Save what's in RAM to a USB Flash drive that contains PSTools and a freeware copy of the program `dd.exe`.**
3. **Shut down the suspect system.**
4. **Reboot to an alternate medium (CD) connected to the suspect system.**
5. **Transfer a bitstream image copy of the suspect system's hard drive to an external hard drive connected to the suspect computer.**

6. Make mirror copies from the original bitstreamed image hard drive copy that you can use for forensic analysis.
7. Create a permanent backup copy of the USB Flash drive contents to preserve RAM evidence.
8. Put the original hard-drive image (created in Step 5) and the USB RAM backup copy (created in Step 7) in a secure location.
9. Copy the contents of the USB drive to the hard drive of an active system for analysis purposes.
10. Examine the bitstream image backup of the suspect system's hard drive — and its RAM — on another computer.

Collecting random-access memory (RAM) dump to a USB Flash Drive

It's often valuable to dump and create a copy of the RAM on a suspect system, because otherwise the evidence disappears when the system is shut down. What becomes unrecoverable is some pretty vital information: Running processes, services, and loaded drivers are the sort of information that may prove relevant to a forensic investigation, particularly if malware processes are running. Some trojans can actually sense and react to changes by removing their physical footprints from the hard drive, but such evidence may still reside in RAM. Additionally, information that may exist in encrypted form on disk sometimes exists in an unencrypted state in RAM.

A RAM dump is investigated with tools that can read hex code. Interpretation of this data requires special skills. From that RAM dump, any relevant data that looks like compelling evidence must be harvested. See the “Getting a Pro to Analyze the Evidence” section later in this chapter for more information.

Keeping track of changes

Any forensic tools you use to investigate the suspect system will produce changes to the file system on its hard drive. Such changes must be documented and presented as evidence of any after-the-fact changes that were introduced. To that end, you can use two Microsoft Sysinternals' programs to record changes to the suspect system state. These changes are output to a log on the USB Flash drive as the evidence is collected.

The two Sysinternals programs in particular we recommend using to record file and Registry activity on the USB drive are called Filemon and Regmon:



- **Filemon** is used to monitor and record file changes in real time and output them to a file on the USB drive. Each file open, read, write or delete operation will be recorded with a time/date stamp. Filemon may be downloaded from the Microsoft Sysinternals Web site:

www.sysinternals.com/Utilities/filemon.html

✓ **Regmon** is used to monitor (and record in real time) the changes that occur to the Windows Registry from attaching a USB device to a system for the first time. Again, it may be downloaded from the Microsoft Sysinternals Web site:

```
www.sysinternals.com/Utilities/Regmon.html
```

This approach works because every time you attach a new USB device to a system for the first time, a Registry entry with the device name, ID, and Serial Number is made under this key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Enum\USB
```

Additional subkeys are created under this key as well, in this format:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Device  
Classes\{CLSID#}\##?#USBSTOR#
```

Regmon can be used to capture these Registry additions to a log file on the USB Flash drive.

This link provides information on USB manufacturer IDs:

```
www.linux-usb.org/usb.ids
```

Installing a copy program on the USB drive

You'll also need to install a copy program on the USB Flash drive to save the volatile physical RAM on the suspect system before shutdown.

A free copy program you can use both to transfer RAM and create a bitstream image of the suspect computer's hard drive is called *dd* for Windows. Actually *dd* is an abbreviated UNIX name that stands for *disk dump* — it's often used to create bitstream image files of media as part of the forensic acquisition process. *dd* is very powerful but its command-line interface is rather complicated. For that reason, here's where we give you the commands required to copy the RAM and hard-drive data from a suspect computer.

dd is made for use on Linux systems. However, there is a version geared specifically for use on Windows platforms called *dd* for Windows. A similar but enhanced program created to collect evidence on Windows platforms is called *dcfldd*. *dcfldd* is so named — *dcfl* + *dd* — because it was independently developed by Nicholas Harbour while employed by the Department of Defense Computer Forensics Lab (DCFL). Either of these programs — *dd* or *dcfldd* — can be used to collect data on Windows systems in the proper format required for forensic analysis.

You can read about and download *dd* or *dcfldd* here:

```
http://dcfldd.sourceforge.net/
```



The Forensics Acquisition Utilities are a “collection of utilities and libraries intended for forensic or forensic-related investigative use in a modern Microsoft Windows environment.” These tools include `dd` for Windows and are discussed at length in the following article:

```
http://users.erols.com/gmgarner/forensics/
```

The Forensics Acquisition Utilities, which include `dd` for Windows, may be downloaded here:

```
http://users.erols.com/gmgarner/forensics/forensic%20acquisition%20utilities-bin-1.0.0.1034%20\(beta1\).zip
```

`dd` for Windows or `dcfldd` must be installed on the USB Flash drive along with PSTools before you can use it for copying data from the suspect system.

Copying the RAM dump

`dd` and `dcfldd` both use a command-line interface, so to issue any copy instructions you must first open a command prompt (Start ⇨ Run ⇨ type **cmd**, and hit Enter). In the following commands, `dd.exe` is an executable file that launches the `dd` program.

You can use the following command to copy active volatile RAM:

```
dd.exe if=\\.\PhysicalMemory of=f:\somepath.dd bs=4k  
conv=noerror
```

In the command just given, `f:\` is the USB Flash drive and `f:\somepath.dd` is the file path for the output file containing the RAM dump on the Flash drive. Make the appropriate substitutions for these variables, particularly if your drive letters differ.



The following `dd` copy command incorporates an MD5 hash function into the copy, to verify the integrity of the output data. This command creates an MD5 hash of the input and verifies the output file to see whether the MD5 hashes match. It also creates a text file called `md5.txt` on the USB Flash drive with a copy of the MD5. Here’s what it looks like:

```
dd.exe if=\\.\PhysicalMemory of=f:\somefile.dd bs=4k  
conv=noerror --md5sum--verifymd5 --md5out=  
f:\md5.txt
```

Now you can use any hash program, such as `Md5Deep`, to verify the integrity of the file at any time. (See Chapter 9 for more information on hashes and other alternative hash-comparison program suggestions.)

```
http://md5deep.sourceforge.net/
```

Using Norton Ghost imaging to create the hard drive image

To make a bitstream image of the suspect system's hard drive, you can use Norton Ghost or `dd`. Using Norton Ghost is a bit easier than using `dd`, so we review that method first.

Norton Ghost (as shown in Figure 10-1) has an option to Create a Ghost Boot Disk with Ghost and other drivers. You can use that option to create a Ghost bootable floppy or CD-ROM for booting the suspect system.




Figure 10-1:
Norton
Ghost menu.

Now that the bootable Norton Ghost CD or floppy is made, insert it into the suspect system and boot.

When you've booted from the CD or floppy, the `-IR` switch is used to specify raw image copying must be done. Specifying raw image copying ensures that the hard drive data is transferred using low-level disk reads and writes that do not alter or format the data in any way — thereby preserving its integrity. The following command should be used to copy the entire disk as is — including the boot sectors and the *slack space* (unwritten portions of clusters that are assigned to files). Use the version that applies to your system:

- ✓ `ghostpe.exe -IR` — For a consumer Ghost version, such as Norton Ghost 2002
- ✓ `ghost.exe -IR` — For a corporate Ghost version

When you run the appropriate Ghost command, a DOS Graphical User Interface pops up that allows you to specify a target and destination disk, and then to start the copy operation.

Refer to the following Symantec document for a full explanation of Norton Ghost sector-copying switches. Here's where to get hold of it:

```
http://service1.symantec.com/SUPPORT/ghost.nsf/docid/200111413481325?Open&src=&docid=19
```

Using dd or dcfldd to create the hard drive image

To use dd to make a bitstream image copy of the suspect system's hard drive, first you have to create a bootable Linux CD.

The following are links to donationware Linux boot CDs that can be downloaded from the Internet:

- ✓ **The Helix Bootable CD:** Helix is available as a free downloadable ISO image. This disk comes with dcfldd, plus other forensic software.

```
www.e-fense.com/helix/
```

- ✓ **FIRE — Forensic and Incident Response Environment Bootable CD:** FIRE is free and it also comes in ISO format so it is ready to be burnt to CD. FIRE provides tools to conduct forensic analysis (including Ethereal) and virus scans on the suspect system. It offers an easy to use, menu-driven user interface.

```
http://fire.dmzs.com/
```

- ✓ **Knoppix:** Knoppix is a distribution of Linux that can also be used to make a bootable CD:

```
www.knoppix.org/
```

- ✓ **The Farmer's Boot CD (FBCD):** FBCD is an excellent payware alternative. It's one of the best Forensic bootable CDs available, but it's also expensive:

```
www.forensicbootcd.com/
```

To make the bitstream image (or duplicate) of the suspect system's hard drive, you'll need to run dd or dcfldd from the Linux boot CD.



Because the Helix Bootable CD comes with dcfldd included no separate download of dd or dcfldd is required if you choose that bootable CD option over the FIRE Bootable CD alternative.

Our example illustrates the commands using dd. You can replace dd with dcfldd in the following commands if you're using dcfldd instead. The commands remain the same except for this minor name swap.



The bitstream image method makes an exact bit-by-bit duplicate of the hard drive that conforms to the exacting specifications required for evidentiary purposes. If an exact bitstream image is not created, then it will not be representative of the crime scene and it may be deemed inadmissible for use as evidence in court.

The following command creates a bitstream image of the suspect system's hard drive, using `dd`:

```
dd.exe if=/dev/hda of=g:\image.dd conv=noerror, sync
```

Here's another command that can be used to create a sector-by-sector copy of the hard drive — but if your aim is to collect legally admissible forensic evidence, the bitstream image copy is the way to go:

```
dd.exe if=/dev/hda of=/dev/hdb conv=noerror, sync
```

The terms used in these commands are defined in the following table:

Command Term	What It Is
<code>if=</code>	The input file or device
<code>/dev/hda</code>	The device or hard drive <i>a</i> (IDE primary master).
<code>of=</code>	The output file or device.
<code>conv=noerror, sync</code>	Instructs <code>dd</code> not to stop on bad-sector errors, and if it finds a bad sector on the input drive to pad the output so the sector's in and out are equal. (If you skip a sector on the input drive, then all physical sectors written to the image are now off by one, and you don't have a true sector-by-sector copy.)

Getting a pro to analyze the evidence

In a court of law, digital evidence is considered to be equivalent to physical evidence. The processing and selective retention of appropriate evidence for forensic analysis is a science in itself that requires precision and expertise as well as knowledge of current laws regarding what constitutes admissible evidence. Though you may have saved the hard drive and RAM data properly, sifting through that data to gather only what you need to make a forensically sound case is . . . well . . . even more difficult.

If you want to pursue litigation, turn an image of the infected system over to the pros, — or even go one step further — by letting them handle the entire evidence gathering and analysis process. To reproduce the criminal state of

the computer or network, the actual harvesting and analysis of the digital data is best done (and best analyzed) by seasoned computer forensic investigators who have proven track records in gathering and presenting meaningful evidence. First of all, their experience enables them to conduct their analysis without tainting the evidence (which may render it inadmissible in court). They also are knowledgeable about the current legal process and know how to filter and extract data suitable for evidence by digging it out of the mountain of collected data.



One such highly skilled investigator is the *Rootkits For Dummies* Forensic Network Advisor, Dave Kleiman. Dave has successfully provided forensic services for years and is highly skilled in evidence acquisition and analysis techniques. Dave knows how to best preserve evidence for use in court. He can perform on-site live acquisitions of your networked systems so you'll experience little or no downtime. Dave's bio is included in the Acknowledgments, detailing his experience and credentials. You can also obtain information on his Forensic Investigation Services by visiting his Web site:

www.davekleiman.com/services.php

A professional forensics investigator can use tools such as EnCase (by Guidance Software) to recover deleted files, e-mails, and file fragments stored in the unallocated space of your hard drive.

If your computer installation is so vital that it cannot afford to experience downtime, then EnCase even allows a forensic specialist to do a live acquisition of the information contained in unallocated disk space (it's normally inaccessible). This live data gathering process requires the client portion of EnCase to be installed and running on the suspect computer at the time of data acquisition. Remember, this setup applies to *live acquisitions* where the computer system remains up and running during the acquisition process. To read up on EnCase, refer to this link:

www.guidancesoftware.com/

Preparing for Recovery

How do you know if you really have a rootkit, before using any of the detection applications described in Chapter 8 and 9? Maybe you don't know for sure. You could have symptoms resulting from a rootkit payload, or just malware in general, and they can be indiscernable. Your symptoms may be very subtle, or extremely obvious and discomfoting. If you know from long use and experience how your computer usually works, and it suddenly starts acting bizarre, you should start by disconnecting from the Internet as we

describe in the next section. Gather what evidence you choose, using the techniques given earlier in this chapter, and then start running the anti-malware and anti-rootkit applications we recommend in Chapter 9. If you get no results that way, then consider using a one of the bootable CDs recommended in Chapter 9. If you haven't created one already, you're going to wish you had (you'll see why soon). If rootkits or other forms of malware are the cause, adopting this approach will soon either put a stop to it or at least make you realize you need more help. If the latter rings true, then solicit the advice of an expert or from the security forums we've mentioned throughout the book and in Chapter 13.

Cutting off network connection before cleaning out the rootkit

Not only do you need to disable your network connections, but if you strongly suspect that you've been hacked, you have to get physical about disconnecting from the Internet. If a rootkit has set up shop in your computer, it may appear to be offline yet still be online because the rootkit's hiding the connected state. *Disable* those network connections, turn off your modem, and unplug the network cable from the back of your computer. *Now* you're disconnected.

Long, long before you wind up having to do these disconnections, be sure that you're prepared for this day. Here's the short list of basic preparations:

- ✔ You've reviewed Chapter 9 and the Appendix.
- ✔ You've created bootable CDs that you can use in the event that your operating system cannot be trusted.
- ✔ You have copies of all applications that can help you with detecting and cleaning malware (including rootkits) on CD.



If you're not a world-class expert user (and most of us aren't) and you detect possible rootkits using your bootable CDs, it's best to call in the cavalry: Find another (uncompromised) computer and use it to visit one of the security Web sites listed in Chapter 13. Make a new topic post in an appropriate forum. Tell them what you've found, what you've done about it, what you suspect, and ask for their help.

Many factors and processes are at work when you're dealing with rootkits and their malware payloads. Removing a rootkit isn't quite as simple as taking out a virus. You really need experienced, expert assistance if you want to remove a rootkit — whether from a standalone computer or from a network server. Of course, if you've really and truly prepared for this outcome,

then you have full, clean backups you can use after following the process presented in Chapter 11. That process of erasing and repairing the hard drive — and then formatting and partitioning it so you can freshly install your operating system — is your last resort if all other efforts fail.

Planning your first reboot after compromise

When you've discovered that your system is compromised, the first order of business is to refer to Chapter 9 (you should first approach detection and removal using the many anti-rootkit tools discussed in that chapter). In most cases, those programs alone are sufficient to properly disinfect your system. After rootkit removal is complete, follow-up anti-malware scans are essential, to remove the rootkit's payload. If, however, using the rootkit-detection-and-removal tools outlined in Chapter 9 have not enabled you to fully detect and remove the rootkit, then refer to the Chapter 9 section on creating and using bootable CDs to get an uncompromised view of the infected OS system. If you have not already created a bootable CD, then you will have to create one (or more) using a clean uncompromised computer at this time. That is why it is so very important to make a bootable CD before you need it, so you have the convenience of using your own clean OS to produce it!

When you have a bootable CD ready, your first reboot after compromise won't be to your own computer's OS; instead you'll be boot up using one of the Windows or Linux CDs you've created. The tools included on the CDs (especially the Windows-based ones) will help you to detect — and possibly remove — the malware and rootkits you find. Then follow the instructions outlined in the upcoming subsections (are also included in the Appendix).

Another reason to use a bootable CD is to double-check your system after you've removed a rootkit using the ant-rootkit tools run from *within* your computer. Booting to a CD in this case provides an additional means of verifying that all rootkit components have been successfully removed. By doing this, you are essentially checking your computer in two different ways (internally during the initial detection and removal process, and externally to confirm removal).

What to save before power-off or reboot

Previous sections of this chapter describe everything you need to save before power-off or reboot. In particular, save your log files, active RAM, and an image of your hard drive on separate media. When those are done, there's one thing left to do, if you've not done so already . . .



Reset your computer's boot sequence *before* you run all those bootable CDs.

How to make your computer boot from a CD

To boot from the PC from a bootable CD, you need to make the CD/DVD-ROM drive as the *first boot device* in the BIOS (Basic Input/Output System). Here's how to get that done:

1. **Restart the computer and keep pressing the Delete key until you see the BIOS setup screen — before Windows loads.**

2. **Change the boot-order priority in such a way that the CD drive (instead of the hard drive) boots first.**

The option that does this job goes by various names, depending on the version and vendor of the BIOS. Usually a section named `BOOT` contains all options related to boot-device priorities.

3. **Use the arrow keys and the + or – keys to navigate through different settings and options in BIOS.**
4. **When you specify the CD/DVD-ROM drive as the first boot device, save the changes and exit from BIOS.**

Don't change any other options in the BIOS unless you know *exactly* what they're for. The BIOS provides loads of functionality, but it doesn't forgive mistakes and can't tell when a change is made in ignorance. If you're not comfortable doing this procedure yourself, please visit one of the security sites listed in Bonus Chapter 2 and ask for help. These two BIOS information Web sites are good resources:

```
http://motherboards.mbarron.net/bios.html
```

```
www.bioscentral.com/
```



Burning an ISO image to CD

To burn ISO image files, you need third-party CD/DVD-burning software such as Nero, NTI CD Maker, Roxio EasyCD Creator, or similar tools. ISO image files should not be directly recorded to a CD as normal data. All CD/DVD burning software will have an option called `Burn Image` or `Record Image to Disk` (or something similar). Use this option to burn the ISO, so the CD will become a *bootable* CD. If you need help on burning ISO images to CD, you can visit this page:

```
www.ntfs.com/iso-burning.htm
```



It is essential that you use a clean computer to create a bootable CD, so having one already made and handy, *before* you need to use it, is your best bet — borrow some advice from the Boy Scouts: *Always be prepared!*

Don't go through this process alone — even if you're sure you can handle it. Please do obtain expert advice and help. If all else fails, you can reboot once more using the hard-drive tools in the Appendix and in Chapters 14 and 15. Be sure to have full, clean backups of your system. Failing even that, be sure to have the Install CD for Windows XP and refer to the instructions provided in Chapter 11. Using another computer, or your telephone, get a second opinion before doing more.

Chapter 11

Preparing for the Worst: Erasing the Hard Drive

In This Chapter...

- ▶ Scrutinizing System Restore (which can't be trusted after a rootkit compromise)
 - ▶ Getting wise to devious malware (and why a simple format and reinstall won't work)
 - ▶ Erasing your hard drive and installing the operating system
-

I imagine that you've done your level best, with expert help, to remove a rootkit and its accompanying payload of malware from your computer . . . and it's a bust. Your system has been corrupted throughout and can't be trusted as far as you're tempted to throw it. If your backups were made before the rootkit moved in — and they were stored on separate media — then congratulations on dodging a bullet. But if not, your existing files are suspect, and you're looking at a fairly drastic scenario just to get up and running again.



If you have not yet been compromised by a rootkit, now is the time to begin making regular backups of your important data and files. Using the Acronis True Image software (included on your *Rootkits For Dummies* DART CD), you can back up your entire hard drive with ISO images saved on CD/DVD. Or you can use the Replicator (also included on the DART CD) to do the same thing without using ISO images. Either way is good, although the Replicator is free for personal and noncommercial use. Keep these regular backup CD/DVDs dated properly and stored in a secure place, preferably locked up in a safe if you have one. (You'll find instructions for using these applications on your DART CD, in the Appendix, and in Bonus Chapter 2.)

Don't Trust System Restore After Rootkit Compromise

Unfortunately, neither your data nor any of your software and files can be trusted after a rootkit compromise. Rootkits work from outside your operating

system; they can corrupt anything it can do, including System Restore. They can falsify or change the dates of the Restore points, as well as the content and data recorded. Even if System Restore has not been corrupted, it can stupidly save both the rootkit and its payload of malware to a special folder that the rootkit reloads every time the system is started or the malware is renamed and deleted.

Rootkits can hide on bad sectors of your hard drive and in the boot sectors. Bad sectors are places where logic errors occur on the disk itself. This can be caused by a physical imperfection or a wrongly magnetized spot. Encountering them could result in system crashes and BSODs (Blue Screens of Death), so your hard drive's firmware maps these sectors in order to avoid them. Rootkits thus have places to hide on your hard drive that are essentially outside the operating system's environment — untouchable by it, yet still at hand.



Firmware is the software that actually resides on hardware components (such as a hard drive in this case) and runs the hardware on your computer. Firmware is usually supplied by the manufacturer. The firmware for your hard drive exists separately from the operating system, yet communicates with it. For example, firmware maps the surface of the hard drive, letting Windows know where files can be stored. It does not allow Windows to use damaged or error-ridden parts of the disk, known as *bad sectors*. A rootkit also exists separately from your operating system and the firmware for your hard drive. Rootkits exist that can copy themselves to bad sectors on the hard drive.

All disks, from floppies to CD/DVDs and hard drives have a boot sector, placed there by the format utilities used on them. These usually contain information about the disk and small programs that provide error messages if the disk is used incorrectly — so you can't use a data disk as a boot disk. Hard drives have another boot sector called the *master boot record* (MBR), which contains instructions for loading your operating system at boot or at startup. Tempting target, isn't it? Well, at least one persistent, kernel-level rootkit can reside in the MBR.



A small number of rootkits can hide in unusual places such as on the BIOS, motherboard, video-card EEPROM (programmable chip), or in alternate data streams. If all other methods fail to eradicate the rootkit — including the procedures provided in this book (yikes) — then it's likely the rootkit is residing in one of those places. If so, it's imperative that you seek expert and experienced help. *Do not connect the infected machine to the Internet. Keep it completely offline until it's clean.*

If your computer has been compromised by a rootkit and its payload, we recommend that you disable System Restore to prevent re-infection by that route. To do so, follow these steps:



You have to be logged on as an Administrator or the System Restore tab will not appear.

1. Click **Start**, right-click the **My Computer** icon, and then choose **Properties**.
2. Click the **System Restore** tab (see Figure 11-1) and check either **Turn off System Restore** or **Turn off System Restore on All Drives**.




Figure 11-1:
The System
Restore tab.

3. Click **Apply**.

Zap. When you turn off System Restore, you delete all previous restore points automatically.

4. Click **Yes**, and then click **OK**.

System Restore is disabled.

To re-enable System Restore, follow these same directions — except *uncheck* the box for Step 2. (Chapter 4 has more about restore points.)



When a Simple Format and Reinstall Won't Work

For years, experts have been claiming that you can delete rootkits from your system by formatting the hard drive and then reinstalling your operating

system. Actually this subject has much more to it than at first meets the eye. Most wiping, erasing, formatting, and partitioning tools will *not* overwrite logical bad sectors on the disk — and that means some rootkits can hide there. To get a handle on this devious tactic, read on.



Bad sectors are physical imperfections and logic errors that exist on the surface of the hard disk itself. A logic error is usually produced by an improperly magnetized area. The write heads sometimes make mistakes, often due to factors beyond their control. Hard drives rarely go out the factory door without some imperfections on the disks. In the old days, hard drives had to list all of their imperfections before you could load and use them. Today, with our more intelligent hard drives that use S.M.A.R.T. (Self Monitoring And Reporting Technology), the disk firmware provided by the manufacturer works outside the operating system. The firmware detects and maps where the bad sectors are on the disk surfaces, excluding them from what your operating system and disk utilities see. Just like a rootkit, eh?

Not only do you have a lying and cheating operating system if you have a rootkit, but your hard drive isn't giving you the complete story — even when it doesn't have a rootkit. It's just a bit too casual about reporting bad sectors. When you check your hard drives and they're given a clean slate, numerous imperfections and logic errors still exist, mapped by the hard drive's *firmware* — which tells the operating system and disk utilities to overlook them.

No wonder that rootkits can hide in bad sectors on your disk. They take advantage of this flaw, provided so generously by S.M.A.R.T. Bad sectors are created most often by crashes of the read and write heads as they move over the disks. A submicroscopic gap exists between the head and the disk. Air pressure keeps them apart but dust, smoke, high humidity, sudden power failures, physical shock, corrosion, and many other horrors can cause the heads to crash onto the disk surface, causing scratches or logic errors. Scratches are physical bad sectors. *Logical* bad sectors are the ones where the read-write heads have crashed without physically harming the surface of the disk, but have magnetized an area improperly; data is lost, producing logic errors. Rootkits can write themselves into those areas — or they may map themselves as bad sectors.

For more information on S.M.A.R.T, see

http://en.wikipedia.org/wiki/Self-Monitoring%2C_Analysis%2C_and_Reporting_Technology

Here's where to find a refresher on how hard drives work:

http://en.wikipedia.org/wiki/Hard_disk

So, what can we do about those logical bad sectors? Obviously now, just erasing, formatting and partitioning does not fix them. We need something that can work without needing your operating system and isn't influenced by the hard drive's firmware.

We found the answer with Steve Gibson of Gibson Research Corporation — specifically with his SpinRite hard drive repair and data-recovery software. SpinRite 6.0 can repair and restore logical bad sectors on a hard drive even after it has been wiped or erased. The software has its own operating system. You need to create a bootable floppy, CD, or DVD with it — when your operating system is clean — for use on your hard drive if you ever need it. Simply put the floppy, CD, or DVD in the appropriate drive and then reboot or restart the computer. Follow the on-screen instructions.

You can get SpinRite and more information about it at the following URL:

www.grc.com/sr/spinrite.htm



Before you start using SpinRite, you need to use a freeware utility known as Eraser to wipe your hard drive completely. Instructions on how to use Eraser are provided in the next section and in more detail in Bonus Chapter 2. You can obtain a copy of Eraser at www.heidi.ie/eraser/default.php. Another freeware application known as Active@ Kill Disk for Hard Drives by Lsoft Technologies Inc. can be used to accomplish both tasks of overwriting logical bad sectors and physically erasing the entire hard drive. All data on the hard disk will be unrecoverable after using it. You can obtain it at www.killdisk.com/. More information about the features for Active@ Kill Disk for Hard Drives can be found at www.killdisk.com/version.htm. Active@ Kill Disk for Hard Drives has a professional edition for commercial uses. Note that you *must* have a floppy disk drive to use Active@ Kill Disk for Hard Drives.

Erasing Your Hard Drive and Installing the Operating System

The purpose for erasing your hard drive is to remove all traces of the rootkit payload of malware, corrupted, altered and modified files and data. This procedure is for use on computers that are so badly infected that no other procedure will get them clean. If there is any way to clean and recover the system *without* using this procedure, use it. If the rootkit infecting your computer is the kind that hides in the logical “bad sectors” on the hard drive, then you'll need to truly erase the hard drive to get rid of the rootkit infection. (It may be good practice to assume this is the case anyway, just to be on the safe side.)



You can't use these procedures to do a sudden erase-and-install on your hard drive at the last minute and expect them to work right. With a rootkit, and without the preparations or the Windows XP Install CD, you would have an expensive doorstop or paperweight. Repair technicians usually have the tools and equipment to fix these problems. If you require such emergency measures but haven't made the preparations listed in the next section, then it's time to bite the bullet: We strongly advise you to take your computer in to the repair shop to let the pros diagnose and repair it.

What you need before you begin this procedure

Here's a handy list of items and procedures you need to take care of *long before* erasing your hard drive becomes imperative. These preparations are indispensable to renewing your computer properly.



Do your best to acquire the following things when your computer is (a) free of all malware and (b) stable:

- ✓ **Full backups** of all the data, files, and software you want to keep.
- ✓ **CDs for your computer hardware devices.** Very important.
- ✓ **Device drivers:** Download the latest versions from the Internet, then copy the drivers on your computer to a CD. That way, after you install your operating system again, you'll have all your drivers up to date. Check the Web sites of the manufacturers for updates.
- ✓ **Motherboard:** Go to the manufacturer's Web site and download the latest drivers for your chipset, and any utilities you use for your version of Windows. You can discover what your motherboard is by checking either the documentation provided with your computer or by Googling "motherboard + *your computer's make and model.*"
- ✓ **Graphics Card:** Download the drivers from either the manufacturer's or the graphics processor's Web sites.
- ✓ **Soundcard:** For separate soundcards, get the full version rather than the update one as this will be a full installation driver set. If you're using the onboard sound system, this is usually included with the motherboard drivers.
- ✓ **Other Components:** Get the latest drivers for your components (see below) and your version of Windows from the manufacturers' Web sites. If these are unavailable, search for them via Google or at a driver-archive Web site such as www.hardwarehell.com/drivers.shtml or www.windrivers.com

- ✓ **CDs and copies of all the security and utility software you use.**
- ✓ Eraser and SpinRite bootable CDs. If you have a floppy drive you can use Active@ Kill Disk for Hard Drives instead.
- ✓ **The Windows XP Installation CD.** Note that you don't want the Recovery CD.

Fortunately, finding the specific information for your soundcard, video card, CD-ROM, and such is pretty straightforward:

1. **Click Start and choose All Programs ⇨ Accessories ⇨ System Tools⇨ System Information.**
2. **Click the plus sign (+) next to Components in the left pane.**
3. **Click the respective heading for each:**
 - **Soundcard:** Choose Multimedia ⇨ Sound Device. The information you need will appear in the right pane.
 - **Video Card:** Choose Display. The information appears in the right pane.
 - **CD-ROM:** Choose Multimedia ⇨ CD-ROM.

Most devices can be found using this method in Windows XP Professional.

Keep all these materials labeled, dated, and stored safely. In fact, it's not a bad policy to keep them in a safe, or a safety deposit box, or a locked drawer in a separate location. Once a week, do incremental updates; each month, do a full update using new discs.



To ease some of the burden of doing the full erase-and-install procedure, you can create a special group of folders beforehand and use them to keep copies of your essential security programs and files. Then, once or twice a month, review it to ensure that you have the tools you want — and burn what you have to CD so you have a complete set of tools someplace other than your hard drive. Keep those CDs with your “erase and install” set, and the work of a new install will be a breeze. Just be sure to have all files and applications unzipped or unpacked and ready to go.

Erasing, partitioning, and formatting

First, before we get started you need to change the boot order in the BIOS of your infected computer. Here's what that looks like:

1. **Restart your computer and press one of these keys: Del, Esc, F1, or F2.**

Different computers use one of these buttons to open it. Watch for a prompt as the computer powers up.

2. If your computer asks for a password before running the BIOS setup program and you don't know the password, take one of two actions:

Take your computer to the repair shop and explain your problem. The ways found to bypass the BIOS password require you either to open the case (thus voiding your warranty) or involve illegal password-recovery techniques.

Consult your computer manufacturer to determine how to open the BIOS safely and legally.

When you get the BIOS setup program running, move to Step 3.

3. Set the BIOS to boot from the CD-ROM as the First Boot Device, with the hard drive as the Second Boot Device.

4. Save and Exit from BIOS Setup.

Be sure to disable System Restore as described earlier in this chapter.

5. With the BIOS ready, create bootable CDs with Eraser and SpinRite.

These are the tools you use to (respectively) erase your drives and kill off any lurking rootkits. Instructions for creating a SpinRite bootable CD are in the Appendix and in Bonus Chapter 2; here are the steps for creating an Eraser bootable CD:

6. Install Eraser onto a clean system; then use it as follows:

- a. Put a blank, writable CD in the drive.
- b. Click Start and choose All Programs ⇨ Eraser ⇨ Create Boot Nuke Disk.
- c. Click OK.

The Boot Nuke Disk is created.

7. Bring the Boot Nuke CD to your infected computer. Place it in the drive and reboot.

All drives will be erased, along with the operating system.

8. Place the SpinRite bootable CD in the drive of the erased computer and reboot.

This ensures that all recoverable bad sectors are repaired; any rootkits still residing in them will be erased. SpinRite has its own operating system — which is not Windows — so it will not be affected by rootkits remaining in bad sectors on your computer.





Rootkits are platform-specific. This means that a rootkit designed for Windows XP will likely not be able to infect the FreeBSD-based SpinRite operating system. To learn more about FreeBSD, go to <http://en.wikipedia.org/wiki/FreeBSD>.

Installing Windows XP

Now that we have wiped the hard drive squeaky clean, removing all files including the infected ones, it's time to install Windows XP. The following instructions assume that your previously infected computer is not connected to the Internet or to a network in any way. Be sure that Internet and network connection plugs are removed to prevent your unpatched operating system from becoming infected by worms and other forms of malware.

Partitioning and Formatting the Hard Drive

Before we install Windows XP (these instructions work for Professional or Home Edition), we need to create one or more partitions on the hard drive and format them by following these steps:

1. Start the computer and place the Windows XP Install CD in the drive.

As mentioned earlier in the chapter, be sure you've set your BIOS to boot from the CD first — *before* you follow these steps. Otherwise your computer ignores the CD and tries to boot from the hard drive (which, of course, is empty).

A message will appear, instructing you to Press any key to boot from CD-ROM.

2. Press any key on the keyboard to boot from the CD.

You see an option to press F6 at the bottom of the screen if you need to install a SCSI or RAID controller; ignore it unless you know you need to. If you do have SCSI or RAID controllers, you can get more info at

www.theelderageek.com/clean_installation_of_windows_xp.htm

You're asked whether you want to install Windows XP.

3. Press Enter to get started.

The EULA appears.



4. **Scroll through the EULA and press F8 to accept it when you get to the bottom.**

After Windows XP installs some setup files, next up are the available partitions where you can install XP.

5. **Pick the partition you want and press Enter (choose the primary drive partition for XP).**
6. **If you have done the full erase of the hard drive, there should only be one unpartitioned space, so press C to create a new partition in space that has not been partitioned.**
7. **Type in the size in megabytes (MB) and then press Enter (or simply press Enter to use the entire drive as your partition).**



8. **In the next screen, use the arrow keys to choose the formatting option you want (NTFS or FAT32) and press the F key to begin.**

NTFS is the way to go if you enjoy your security and want a stable operating environment.

9. **Choose the type of format you want.**

We would advise you choose “Format the system using the NTFS file system” because you are starting over with a completely clean hard drive. The full format will check for and map the physical bad sectors on the hard drive. The options with <Quick> beside them do not check for bad sectors.

When the format is done, installation resumes. Eventually the computer reboots.



When the computer reboots, *do not press any key* when the *Press any key to boot from CD* message comes up. Otherwise you have to start the whole process again.

10. **As the installation proceeds, supply appropriate responses as you see fit to complete the setup.**

After another reboot you’ll be at the Windows XP Activation screen.

11. **You can activate now or any time in the next 30 days. (Registration is optional.)**

12. **Choose a username.**

The Windows XP desktop will appear.

If you’re using Acronis True Image, install it and make an image of your pristine and pure drive.

Are we really done yet?



After you install . . .

Before you install your other software and files, please run several of the rootkit detection applications we have provided, just in case. If all indications are clear and clean then start installing your remaining software and drivers.



When you get your computer working with all proper drivers and applications installed, make another image of your drive so that if you should get infected again and cannot get cleaned up, instead of going through this whole process, you can fully erase the hard drive, install Windows XP, and then apply your drive image. All of the drivers and your favorite software will be installed and ready. It's a worksaver!

If you have any problems with shutdowns such as BSODs or lockups after installation, please refer to the following article, “The Role of the F5 Key and Shutdown Problems,” at this address:

www.theelderageek.com/clean_installation_of_windows_xp.htm

. . . And beyond

If you haven't already, consider checking out the chapters at the start of this book (Chapters 1 to 6). By doing so you can glean many tips and techniques to help strengthen your computer against malware and rootkits. In those chapters we show you how to configure a “security geek” computer that's armed to the teeth and ready for almost anything the Internet can throw at you.

For more advice or help from the authors — and from some real malware and rootkit experts — visit this address:

www.castlecops.com/f233-Rootkit_Revelations.html

Part V

The Part of Tens

The 5th Wave

By Rich Tennant



“Oh, Arthur is very careful about security on the Web. He never goes online in the same room on consecutive days.”

In this part . . .

It's the traditional *For Dummies* Part of Tens. In these short-but-sweet chapters, we offer some handy resources: a short-list of major rootkits, and reinforcements when you may need them on the Internet battlefields.

We put ten rootkits in the lineup and summarize their antisocial behaviors — exploring the various techniques they employ to achieve their stealth. Next, if you do find yourself the victim of a malware or rootkit attack, don't feel as if you're left alone to sink or swim. The help is out there — and we show you where to get it: Ten highly regarded online security forums where skilled volunteers can assist you in system assessment and malware removal. (In two bonus chapters on the CD, we add nearly thirty applications with instructions to help you fight and survive the onslaught of rootkits, and to complement the DART-CD of tools.)

For Dummies traditions cannot be denied.

Ten rootkits described (with one besides),

No longer can they hide.

Ten Web sites securely provide

Safe haven for all who come by.

One score and ten programs do guide

Keeping your rootkit security high!

Chapter 12

Ten (Plus One) Rootkits and Their Behaviors

In This Chapter

- ▶ HackerDefender
 - ▶ NTFSsider
 - ▶ Elite Toolbar
 - ▶ Apropos rootkit
 - ▶ FU
 - ▶ FUTo
 - ▶ MyFip
 - ▶ eEye BootRoot
 - ▶ FanBot
 - ▶ pe386
 - ▶ Shadow Walker
-

Some people have said that rootkits by themselves are innocuous. We beg to disagree; their cloaking abilities make them a form of cyber-piracy. As with the term *blackhat hacker*, rootkits have become associated with evil undertakings; if they were once relatively innocent, they are so no longer. The ten (plus one) rootkits presented in this chapter are a small sampling of the most current and widespread forms. As in battle, so it is with computer malware: Know thine enemy. The descriptions collected here can help you get a handle on these malware programs — from the pernicious and notorious species lurking in the wild to the scary-but-tame (so far) laboratory examples being used to improve anti-rootkit defenses.

HackerDefender

HackerDefender is a fairly common rootkit which has been in existence since 2002. It is also called HxDf, and it has several variants. HackerDefender installs a backdoor and downloads a huge malware payload onto the infected computer. Even after the unclocking and removal of HackerDefender, this folder must also be tracked down and removed. As with most backdoor trojans, HackerDefender takes its initial directives from the information contained in an initialization file named `Hxdef100.ini`, which resides in the `Hxdef` folder. As described in Chapter 9, HackerDefender installs a kernel-mode driver called `hxdef100drv.SYS`. The driver allows HackerDefender to achieve kernel-mode functionality by hooking APIs referenced in kernel data structures. The HackerDefender driver is loaded by the HackerDefender service (`hxdef100.exe`); when this process is stopped, APIs are no longer hooked — and the rooted files become visible.

HackerDefender sniffs data transmitted on an already-open and heavily trafficked port (such as the HTTP server port 80 or SMTP port 25) to capture commands from its remote master. Chapter 8 details the vulnerabilities of ports, as well as how backdoor threats often target reserved ports to listen for commands or transmit information — in ways that don't trigger a resident firewall or an IDS block. HackerDefender hides itself thoroughly, tucking away anything that begins with the prefix `hxdef` in cloaked locations on disk, and sometimes opening additional ports that traditional port-viewing application can't see.

Many rootkit-detection programs can detect HackerDefender's hidden service and/or driver:

- ✔ The service and driver files should be visible via IceSword's Port function.
- ✔ Using TCPView or Port Explorer in combination with AntiHookExec (as described in Chapter 9) should turn up those files on-screen as well.
- ✔ Microsoft offers the Malicious Software Removal Tool (MSRT) that silently scans for HackerDefender and several other rootkits every time you download and install updates for Windows.

You're only be notified of MSRT's scan results if malware is detected — so you may not even realize this background scanning is taking place. With the MSRT, no news is good news — but it's reassuring to know that this rigorous scanning process happens each time updates for Windows are installed (yet another reason to download and install Windows updates in a timely manner). You can run the MSRT whenever you like by going to this link:

www.microsoft.com/security/malwareremove/default.msp



NTFShider

This nasty piece of work is a covert file-system driver — a rootkit whose only purpose is to stockpile and hide data and files in bad sector blocks or unused clusters on the hard drive. NTFShider is typical of the rootkits referred to in Chapter 11 — those that require more than simple formatting or partitioning to remove them. So-called “writable bad sectors” must be repaired and overwritten to remove any data that this type of malware deposits there.

You could go to all the trouble of formatting and re-partitioning your hard drive, but NTFShider is a slick customer. If it was lurking there in the first place, it will still be there when you believe it’s gone — unless you use a file-system repair utility such as SpinRite, which is also discussed in Chapter 11, as part of the procedure to restore and renew your hard drive close to its original state. Instructions for using SpinRite can found in Bonus Chapter 2.

Elite Toolbar

A pernicious adware and browser-search hijacker, Elite Toolbar uses rootkit techniques to persist on a user’s computer, making it difficult to remove. It displays unsolicited pop-up advertisements, hijacks the Internet Explorer homepage, and does not provide either a privacy policy or an End User License Agreement (EULA), either of which might otherwise tip off a user to its sleazy nature. It can be installed by the Buddylist Trojan, or come bundled with “free” games such as enBrowser or Snackman, which installs over 30 separate adware and trojan applications.

The primary executable file extracts two DLL files to the same folder and then initiates them. These particular DLL files cause Windows to conceal the Elite Toolbar files and the installation folder.

Also known as SearchMiracle, EM Toolbar, Enternet Media Toolbar, and other such names, Elite Toolbar can cripple your system and drastically slow its performance with a barrage of pop-ups. In addition to this built-in hassle, some of these pop-ups may contain material that you find offensive. The SearchMiracle Elite Toolbar (gotta love those snazzy adware names) has been linked to data theft from the computers on which it’s installed; it adds and removes toolbars from Internet Explorer and causes frequent Blue Screen of Death (BSOD) crashes and system lockups.

A number of commercial products will remove this menace, but you can also seek help from the security sites listed in Chapter 13. You may have more problems than just this one, so it's important to get help from people who know what to do.

Apropos Rootkit

ContextPlus, Inc. has recently announced that it has stopped distribution of its commercial adware program called Apropos. Apropos uses rootkit technology to hide its program components so they'll be more difficult to remove. Since many users have been affected by this rootkit, it will no doubt be a source of problems for some time to come, so we've included it in our rootkit discussion despite its makers' kindly disclaimer.

Apropos is used to spy on and record an infected user's surfing habits. This information is sent back to the Apropos servers and a customized pop-up advertising campaign is generated, based on the information that was transmitted back.

Apropos installs a kernel-mode driver and hooks both user-mode and kernel-mode data structures to achieve its stealth. It installs a spyware service that begins before Windows starts. A hidden spyware folder located in the Program Files folder accompanies its installation. Because of the rootkit, the Apropos spyware program's components are hidden. In addition to using a rootkit, Apropos uses a random renaming scheme: This tactic compounds the difficulty of uninstalling Apropos. Not only is a unique random name assigned when it's initially downloaded from the ContextPlus servers, but the Apropos driver (SYS file) is newly renamed every time an infected user reboots. When program components persistently change names, they're harder to hunt. Automatic scanners rely on consistent definitions that include consistent names, so this tactic makes definitions difficult to develop; it also complicates manual identification and removal. (Interestingly, as discussed in Chapter 9, Rootkit Revealer does the same thing to avoid being successfully blocked by malware.)

Because the Apropos service does not run in Safe mode, its randomly named spyware files, Registry autostarts, and processes are visible in Safe mode. The Apropos hidden spyware service and the driver it loads should all be visible in Safe mode.

BlackLight, Rootkit Revealer, and IceSword can all detect Apropos.

A special Apropos Fix was developed by Swandog46. It's used routinely on the online security forums to effectively remove Apropos. We can hope that the Apropos rootkit will experience a quick demise, now that it's no longer actively distributed — but vigilance is a surer bet.

FU — the Malware That's Also an Insult

The FU rootkit (hint: it has nothing to do with the words “foo” and “bar” that programmers often use as harmless examples when they talk shop) is one of the most sophisticated rootkits out there. It exerts its effects solely through DKOM — for openers, removing itself from the list of active processes and hiding its rootkit activity from the Windows Event Viewer. It performs this bit of magic by upgrading user privileges for the tokens that define security contexts (that is, privilege levels) for specific processes or threads. For example, if a process is run with administrative privileges, its token looks much different from a token associated with the same process being run in a limited user context. In effect, the user-level tokens are selectively disabled to fit the privilege levels of their processes. By replacing disabled tokens with newly added tokens *that have system privileges*, FU fools the operating system into sanctioning an action that would normally trigger an event-log entry. FU accords itself system privileges so its activities can evade Event Viewer notification. A rootkit running with full administrative privileges can override the permissions associated with any object — and files or processes are objects.

FU also uses DKOM to hide kernel device drivers — in fact, it uses DKOM for everything. Thus any rootkit-security program that relies on the detection of API hooks can't find FU — or any other rootkits that employ only DKOM.

FU is an open-source rootkit that was originally a laboratory-designed rootkit made as a proof of concept that has since (unfortunately) been made readily available for download on the Web. FU itself does not hide its own driver file, but FU can be used to hide components belonging to other malware. A modified version of the FU rootkit has been borrowed and incorporated into many threats to achieve stealth characteristics. For example, several dangerous backdoor Bots and the MyFip worm use a variation of the FU kernel-mode driver. By using the FU code, a malware writer is relieved of the burden of having to write their own clean driver code. Even the user-program control functions needed to interface with the driver can be borrowed from pre-existing code. This makes it easy for malware writers to impart stealth characteristics to a threat without having to undertake the difficult coding tasks. The FU trojan dropper (called `fu.exe`) installs the FU driver called `msdirectx.sys`. When FU's kernel-mode driver is installed, it can be used to hide malware processes of other applications.

FU uses sophisticated DKOM techniques to hide malware processes from the process list without or hindering or preventing their execution. BlackLight and IceSword are both able to detect and restore the Process List which causes processes cloaked by DKOM to become visible. A BlackLight scan will not flag the FU driver, however — because it isn't a hidden file. Because the FU driver is not concealed, it lends itself to removal by traditional antivirus and anti-trojan scanners. Rootkit Revealer cannot detect malware processes hidden by FU because FU uses DKOM to alter kernel data structures directly and doesn't derive its stealth from the API-hooking techniques that Rootkit Revealer detects.

FUTo

As discussed in Chapter 7, FUTo is a laboratory-designed rootkit. Spurred on by the fact that FU was able to be detected by some existing rootkit-detection programs, Jamie Butler (the creator of FU) and Peter Silberman were motivated to develop a variation of FU called FUTo. The FUTo authors introduced a new process hiding strategy in FUTo, which successfully outsmarted all current rootkit detection techniques. Both IceSword and BlackLight can detect Direct Kernel Object Manipulation (DKOM) of the process list, a dynamic data structure of active processes which is located in kernel memory. Neither IceSword nor BlackLight can detect FUTo's method of hiding processes using DKOM of the PspCidTable. The PspCidTable is a memory-resident kernel data structure that contains information on both active processes and threads. Not only did FUTo's authors introduce a new method of hiding processes, they also developed a new rootkit-detection program called RAIDE (Rootkit Analysis Identification Elimination) — which can detect FUTo. Currently, RAIDE is only available for beta testing — but it's expected to be released to the general public in the near future.

MyFip

MyFip is a memory-resident worm with a number of variants. It often arrives with a phishing e-mail message that uses an IFRAME exploit (see Table 4-1 in Chapter 4). It drops a trojan component that steals .DOC and .PDF files. It can propagate through network shares by using a list of weak usernames and passwords. When it's logged on, it copies itself to the system and registers DFSVC.EXE as a service called Distributed Link-Tracking Extensions.



Phish are fraudulent e-mail messages sent for the purpose of getting unwary users to visit malicious Web sites to steal their private data such as their PIN numbers and other sensitive information, or to get them to download malware and rootkits. For more information on this illegal practice, go to

<http://en.wikipedia.org/wiki/Phishing>

Any of the MyFip variants may drop another component called MyFip.H, which uses rootkit techniques (such as manipulating kernel data structures) to conceal itself. MyFip.H does not propagate; it needs administrator privileges to proceed. So it installs Registry keys to remain resident, and threads to prevent removal. It steals files from all drives with the following extensions: DOC, PDF, SCH, DWT, DWF, DWG, SCH, MAX, MDB. If it does not find any such files, its process ends but it stays resident, awaiting its prey. The stolen files are sent to a remote server using a TCP port. One of the things that makes MyFip stand out from the crowd is that it uses an unusual multiple-compression technique to fool most e-mail scanners.

Fortunately, MyFip can now be removed by most reputable antivirus applications.

eEye BootRoot

eEye BootRoot was a project presented at the Black Hat USA 2005 Convention by Derek Soeder and Ryan Permech of eEye Digital Security. The purpose was to explore technology that uses boot-time, real-mode code to infiltrate and sabotage the Windows NT-family kernel as it loads.

It acts as a removable-media boot sector that positions itself to run later — as Windows is loading — and then continues the boot sequence from the kernel. Its hooking technique causes a Red Screen of Death (RSOD), a full kernel crash.



If you should ever encounter an RSOD, be sure to check with the experts at the security sites listed in Chapter 13. Or take it immediately to the repair shop.

FanBot

The FanBot Family is a group of e-mail worms that use rootkit techniques to propagate and maintain themselves. Having its own SMTP engine, this worm can copy itself to e-mail messages it sends to addresses it collects from the

Windows Address Book (WAB), or by creating new addresses from ones it has collected before. It also propagates using peer-to-peer (P2P) file shares, using filenames designed to attract people to download them. These activities consume huge amounts of system resources and bandwidth. It creates Registry entries to load itself on system startup, and to disable Internet Connection Sharing (ICS) so users cannot access the Internet. It makes changes to the HOSTS file to deny access to antivirus and security Web sites. FanBot can also disable many processes — including antivirus programs, rootkit scanners, and Windows Task Manager. It can act as a backdoor, opening a random port to allow a malicious hacker access to the computer. The system is opened to other forms of malware attack as well.

If you have any of the symptoms described here, please seek help and advice from a security site such as those listed in Chapter 13 — and (oh, yeah) use a different (uninfected) computer when you do, if at all possible.

pe386

As mentioned in Chapter 7, pe386 is a crafty new rootkit that either newly introduced or revived several interesting techniques. pe386 was initially named after its author, but even though antivirus vendors have since referred to it more formerly as Mailbot.AZ or Agent.Rustick, the pe386 name is what stuck. The pe386 rootkit is hard to detect and remove because it has no running processes and its driver is not only hidden by the rootkit, but it's located in the alternate data stream (ADS) of the `system32` folder. ADS are special storage areas in files that are infrequently used and not normally viewable without using a special program developed for that purpose. Although some antivirus and antitrojan scanners do provide an ADS scanning option, it is rendered useless if the pe386 rootkit is running. If pe386 detects certain popular rootkit detectors are active, it responds by altering its behavior to successfully avoid detection.



Unlike most other kernel mode rootkits, pe386 uses a SYSENTER hook, rather than hooking the usual kernel data structures that most anti-rootkit programs can easily detect. Some anti-rootkit developers have reacted to pe386 rootkit's use of SYSENTER hooking (which is actually an older technique) by implementing SYSENTER hook detection into their programs, while others, like SVV, already included it to begin with. Consequently, pe386 is now detectable by many anti-rootkit utilities — but only GMER and AVG Anti-Rootkit are able to both detect and *remove* it.

Shadow Walker

Shadow Walker is actually a product of the good guys — a proof-of-concept rootkit developed to study new stealth techniques that could potentially be adopted by malware writers in the future. It exists only in the laboratory at this time and was implemented using a modified version of the FU rootkit.

Shadow Walker works by hooking the Virtual Memory Manager (VMM), which enables it to remain invisible to signature-based scanners. Because Shadow Walker exists in memory only, it leaves no files, folders, Registry entries, or other physical evidence on the hard drive— which makes it very difficult to detect. In fact, Shadow Walker was formerly invisible to *all* existing rootkit-detection tools — until a new rootkit detector called RAIDE was developed.

To get a handle on what Shadow Walker can do, a brief review of some points discussed in Chapter 7 is in order:

- ✔ Kernel-level rootkits are implemented through drivers — and those exist in the kernel's virtual memory.
- ✔ Normally AV and anti-trojan scanners depend on signature databases for recognition — a scanner can detect kernel rootkit drivers if the rootkit's signature is already part of the scanner's database — and the rootkit is not running (but that's a big if!).
- ✔ Non-persistent rootkits exist in memory only and are harder to detect because they leave no physical traces of their existence on the hard drive.
- ✔ Because it is very difficult for a rootkit to mask itself in memory, rootkit coders normally forego that extra precautionary step and do not attempt to hide the rootkit. As a result, in-memory rootkits are normally detectable by traditional memory scanners

Unlike most in-memory rootkits, Shadow Walker can conceal its driver (and the modifications it makes) in kernel memory. It achieves its stealth by creating its own page-fault handler to control accesses to virtual memory. If memory accesses are determined to be of kernel-mode origin, then hooked rootkit pages may be swapped in and executed, in the place of bona-fide program instructions. Memory accesses for executing user programs are diverted to the real operating system's page-fault handler. Read-and-write memory operations are also handled appropriately to conceal evidence of the rootkit.

Most programs are too large to exist in the space allocated to them in physical memory or RAM. Instead, they are broken up into blocks, and only a few of these blocks reside in RAM, while the rest of the program resides in virtual memory on your hard drive. When an executing program requests code that is not in physical memory, a *page fault* is generated and the requested page is swapped in from virtual memory on disk. A page that has not been used recently will be swapped out of RAM to allow this transfer to occur.

The *page-fault handler* controls this swapping in and out of program code blocks by issuing I/O requests during program execution. A page table is used to look up the physical page that a virtual page corresponds to — and to see whether it exists in physical memory. Shadow Walker takes advantage of memory swapping to get its rootkit code executed. It adjusts the information in the page table so every page that's accessed appears to the system as if it isn't present in RAM. That's how Shadow Walker diverts all memory accesses to its own page-fault handler to determine the next course of action. It treats program-execution accesses differently from read-and-write-memory accesses. Execution accesses are also handled differently depending on whether the access is in user mode or kernel mode. If the access originates from the kernel and the requested page is hooked, then rootkit code is swapped into memory for execution. If the memory access is a read or write operation, this most likely means another program (such as a memory scanner) has initiated the action. In this case, the page is mapped to a phony legitimate page that won't cause scanners or detection programs to notice any foul play. This tactic effectively covers all bases, and allows the rootkit code to remain hidden.

In summary, Shadow Walker achieves its stealth by

- ✔ Implementing its own page-fault handler to take care of page swaps during code execution, and passing unhooked execution accesses to the normal operating system's page-fault handler
- ✔ Distinguishing between execute and read/write memory accesses
- ✔ Diverting hooked execution accesses to rootkit pages
- ✔ Diverting read/write memory accesses (such as those initiated by scanners) to fake "normal" pages that appear legitimate to memory scanners

Shadow Walker is a deliberately devious nightmare rootkit conjured up in the lab to prepare us for things to come. Whatever the next generation of malware looks like, you can bet somebody will be thinking about new ways to do a rootkit's dirty work.

Chapter 13

Ten (Plus Two) Security Sites That Can Help You

In This Chapter

- ▶ Aumha
- ▶ BleepingComputer
- ▶ CastleCops Security Professionals
- ▶ Geeks to Go
- ▶ Gladiator Security Forums
- ▶ Malware Removal
- ▶ Sysinternals Forum (sponsor of Rootkit Revealer Forum)
- ▶ SpywareInfo
- ▶ SpywareWarrior
- ▶ Tech Support Guy Forums

The security communities on the Internet have a slew of sites that can help you with cleaning malware, hardening your security, and meeting people who really and truly care about everyone's online well-being. We would be happy to list them all, but that would need another book unto itself. The sites we do list here have knowledgeable, competent, and experienced people who can help you with your computer security needs, especially for providing advice about all kinds of malware, including rootkits. Many of the staff and members in good standing at these sites have dedicated their free time, abilities, and expertise free of charge to help you with your hardware, software, and security. Most sites gratefully accept donations to help with the costs of providing these free services.



Microsoft has recognized the efforts of many volunteers online with the MVP Award (Microsoft Most Valuable Professionals). The award is given to those who have made outstanding voluntary contributions to the online Windows communities. Lots of MVPs are on staff, or are members of these security sites. Here's where to find more information about them:

mvp.support.microsoft.com

Note that below you will see many references to HijackThis forums. HijackThis is a freeware program developed by Soeperman Industries (Merijn.org) that helps experts detect and clean browser hijackers and spyware from users' computers.

Aumha

www.aumha.org

Aumha is a well-frequented forum providing HijackThis analysis, assistance with various operating system platforms, and numerous Self Help guides. (HijackThis is a Registry scanner that experts use to scan for malware infections.) The forum offers the popular Parasite Fight Quick Fix Protocol, parasite-fighting recipes, and a lot of sound advice. Not only does Aumha provide an abundance of very helpful information but you'll enjoy reading it, thanks to Aumha owner James A. Eshelman's great sense of humor. Among the handy features is My Favorite Freeware, an area where forum staff members and readers are encouraged to recommend their favorite downloads — and we learned about a favorite text editor (EditPad Lite) from a recommendation found there. Aumha's Top 20 Most Visited Pages has a list of articles on a variety of computer subjects, from Stop messages to batch-file commands to recovering deleted files. The staff is absolutely top-notch; if they can't get to the "root" of your problem (no pun intended), nobody can.

Bleeping Computer

www.BleepingComputer.com

Bleeping Computer began as a forum dedicated to user education through its Spyware Removal Guides & Self-Help and Reading Room offerings. These tutorials and guides span the gamut. Some acquaint users with basic computer concepts, some explain more advanced computer subjects, and others

address the usage of helpful programs or tools. The Spyware Removal Guides, for example, enable users to remove some of the most tenacious infections on the Internet. The guides present their instructions so simply that removing spyware is a relatively straightforward and painless process, even for computer newbies.

Over time, Bleeping Computer has expanded and improved its services. It now has a busy HijackThis forum — and a HijackThis training school for those who want to learn how to assist forum visitors with HJT log analysis.

Bleeping Computer provides support for a variety of operating systems platforms — including Windows 95/98/ME, 2000/2003, XP, Linux, Unix, Apple Mac OS, and a discussion forum on Windows Vista that began with the beta version. Bleeping Computer also maintains a very complete Startup Programs Database; you can use it to research the programs that automatically load and run every time you start your computer. Our *Rootkits For Dummies* Technical Editor, Lawrence Abrams, owns the Bleeping Computer Web site — and he is very hands-on in administering its daily operation. Bleeping Computer boasts an expert staff, and within a very short time it has become one the most active security forums on the Web.

CastleCops Security Professionals

www.castlecops.com

CastleCops started just over four years ago; back then it was called ComputerCops (the name was later changed). Developing rapidly from a new and innovative security Web site, it has become a lot bigger but it still keeps innovating and growing. Started by owner Paul Laudanski (who is also a fully hands-on administrator) and his wife Robin, CastleCops continues to grow and blossom, attracting the most creative, intelligent, and productive experts in security and related fields. All staff members are dedicated volunteers who are guided by the CastleCops vision, which you can find at

www.castlecops.com/t63382-Our_Vision.html

A great many of the staff are MVPs due to their passion for their expertise, including both authors, the tech editor, media advisor, and many of the research team who worked on this book. The book writing and research team have created a group of forums called “Rootkit Revelations” (named for Larry’s first published article on rootkits) where users can get help, cleanings, and advice from rootkit experts. The rootkit forums also plan to open a Rootkit Academy to teach people how to remove rootkits safely.

www.castlecops.com/f233-Rootkit_Revelations.html

CastleCops has a thriving HijackThis forum, one of more than a hundred forums in over 30 sections. To aid the many posters to the HijackThis Forum seeking help for malware infections on their computers, the CastleCops Staff created the “Malware Removal and Prevention” procedure:

<http://wiki.castlecops.com/MRP>

CastleCops has forums for languages other than English, a German-language subsidiary site (<http://de.castlecops.com>), and its own Wiki portal at

<http://wiki.castlecops.com>

Along with Sunbelt Software Ltd., CastleCops has spearheaded the “Fried Phish” Phishing Incident Reporting and Termination (PIRT) Squad to combat fraudulent phishing Web sites and e-mails. You can contact the PIRT Squad at:

www.castlecops.com/pirt

Following on the success of PIRT, CastleCops has created a new way to fight malware with the Malware Incident Reporting and Takedown Squad, or MIRT:

www.castlecops.com/c55-MIRT.html

Geeks to Go

www.geekstogo.com/

GeeksToGo.com is a popular computer technical support forum that offers visitors assistance on all versions of the Windows operating systems from Windows 95 on up, as well as a discussion forum on Vista. There is also Apple and Linux support.

Because we’re very big on user education, one of our favorite sections is the Guides and Tutorials area where users can find information on a variety of topics so they can help themselves. These guides were created in response to the most frequently asked user questions. Here’s a sampling of topics covered: Easy XP Tweaks, Windows XP Blue Screen of Death STOP Codes, and Free Antivirus and Anti-spyware Software. As is typical for online security forums, the Malware Removal – HijackThis Logs forum is the busiest area at

GeeksToGo; you can count on getting highly skilled staff attention for your HijackThis log within a reasonable response time. GeeksToGo also maintains Geek U, a school for training new malware fighters.

Gladiator Security Forum

<http://gladiator-antivirus.com>

Gladiator Security started over four years ago, at first as the flagship product of a new anti-trojan and antivirus software company of the same name. The software author abandoned the project, but his backer, Udo Laumann, did not, and neither did the many good and talented people at the Gladiator Forums. Thus we have the Gladiator Security Forums continuing to grow in the fight against malware and rootkits.

Gladiator not only helps you to clean malware off your computer, they encourage you to understand it — providing excellent information and discussion forums. There you can learn more by asking questions of their experts and joining the quest for Internet enlightenment. Many of their staff and members are active at the other forums we've listed here as well.

Malware Removal

www.malwareremoval.com

The Malware Removal University exists at this site. New student volunteers are enrolled in courses designed and guided by experts to remove malware from computers. Upon graduating, these newly minted malware experts continue the work at other security sites all over the Internet.

The malware-removal forums provide a means by which the students can gain hands-on experience — while guided by expert teachers — in real disinfection methods and techniques. No student fees or tuition are charged. Students need only time, a willingness to learn, and a firm dedication to complete the courses.

Started and owned by ChrisRLG, with assistance by Nellie2 (both are ex-teachers from the TomCoyote Forums), this site continues to be one of the leaders in the online security communities. Here's its URL:

www.malwareremoval.com/myinfo.html

Microsoft Newsgroups

www.microsoft.com/communities/newsgroups/default.aspx

While they're not sites per se, we would be remiss if we failed to mention the Microsoft Newsgroups. Many of the experts that volunteer at other online technical support forums also offer solutions and advice on the Microsoft Newsgroups, which address all Microsoft products. Many of the Microsoft Most Valuable Professionals (MVPs) share their knowledge and expertise in these discussion groups.

Sysinternals Forum (Sponsor of Rootkit Revealer Forum)

www.sysinternals.com/Forum

Sysinternals is the home forum of Mark Russonovich, discoverer of the Sony rootkit and developer of Rootkit Revealer. Rootkit Revealer can be downloaded at the Sysinternals Web site, and you can also receive help with interpretation of Rootkit Revealer logs on the Rootkit Revealer Forum. Although the forums say to "Post logs for input from the community," RKR Forum Moderator namrehto and Senior Member SpannerITWks are well seasoned in Rootkit Revealer log analysis and are there to provide knowledgeable advice. Besides offering Rootkit Revealer, Sysinternals provides free downloads of some of the best Windows system-analysis tools available (such as Process Explorer, Autoruns, Filemon, Regmon, and PSTools) and sponsors dedicated support forums to provide assistance with most Sysinternals programs. Here you can receive answers to your questions and assistance in interpreting your scan reports.

SpywareInfo

www.SpywareInfo.com

SpywareInfo has become a large and busy security site dedicated as it is to removing malware of all kinds from people's computers. As its name implies, if you need information and help from experts with spyware and malware,

this is the place to get it. The forum is easy to understand and use, but do please read their FAQ at

<http://forums.spywareinfo.com/index.php?showtopic=227>

Founded, and developed up until recently by Mike Healan, SpywareInfo continues to provide its excellent services with the backing, assistance, and donations of members of the security community.

SpywareWarrior

www.SpywareWarrior.com

SpywareWarrior is famous for its Rogue/Suspect Anti-Spyware Products & Web Sites — in particular, this one:

www.spywarewarrior.com/rogue_anti-spyware.htm

This gem is a compendious and regularly updated list of fraudulent and spurious anti-spyware applications and sites. The list helps consumers to make wiser choices before buying anti-spyware products online — and to remove false ones from their computers. If you're having problems with spyware and intrusive adware, they can really help you here.

SpywareWarrior was founded by — and continues to be developed by — Suzi Turner and Eric Howes. Both are big names and bona fide experts in the security field. They provide a tremendous amount of concise, clear information on all things spyware. Their staff, too, is second to none, providing time, effort, and expertise that meet the highest standards.

Tech Support Guy Forum

www.techguy.org

Tech Support Guy Forum was one of the first online forums to offer free technical advice, and it is celebrating its tenth anniversary this year. The “Tech Support Guy,” aka Mike Cermak, runs a manageable but very active forum where you'll receive highly qualified assistance. Though very busy, the forum sports a streamlined appearance, so you'll have little trouble locating where to go to get what you need. Assistance is provided for users running

Windows 95, 98, NT, ME, 2000, XP, UNIX/Linux, and Apple Macintosh. Forum visitors can post HijackThis logs in the *Security* section, which is usually the busiest area of any online security forum, and it is no different here. If you do post a HijackThis log, not only will you receive expert advice, you'll receive it extremely quickly. Most forum staff members have been dispensing advice for a long time, and they really know the ropes. You'll find yourself in extremely capable hands at the Tech Support Guy Forum.

Tom Coyote Security Forum

www.TomCoyote.com/

Tom Coyote is held to the same high standards as GeeksToGo, as both are owned by Tony Blair and Tom Wilson. It also has an active HijackThis Forum. Tom Coyote was one of the first sites to offer a training program for malware fighters via the Tom Coyote Classroom, which is still flourishing today. As far back as 1995, Tom Coyote Wilson was one of the first people to pioneer security Web sites on the Internet. Working alongside greats such as Steve Gibson (SpinRite and Gibson Research), and Nicholas Stark (Lavasoftware and Ad-Aware) in the early days, he promoted safe and secure computing on the Internet for everyone. He's still doing it, and if you need help with your computer and Internet problems, visit his Web site.

Appendix

About the CD

In This Appendix

- ▶ System requirements, installing, and troubleshooting
 - ▶ Anti-malware utilities and scanners
 - ▶ Backup and imaging applications
 - ▶ Rootkit-detection-and-removal applications
 - ▶ Password protectors and generators
-

Throughout the book we've been gushing over our beloved DART CD (or Dummies Anti-Rootkit Toolkit) and all of the fantastic utilities we got to put on it; this appendix describes what you need to run the software on the DART CD, how to install it, and other technical stuff we thought you might need to know.

System Requirements

The entire course of this book assumes you're using Windows XP Home or Professional, with a few references to Windows 2000 or Server 2003. Some of the applications provided here will run under all recent Windows versions. Each entry has a notation that spells out its specific requirements and compatibilities as provided by the developers. For more detailed instructions, please refer to the online documentation provided by the vendors of these software programs.

Make sure your computer meets the minimum system requirements shown in the following list. If your computer doesn't match 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 in the root directory of the CD-ROM).

Some applications may permit lesser requirements than those listed here.

- ✔ A PC with a Pentium 3 or faster processor; 450 MHz or higher CPU speed
- ✔ Microsoft Windows XP Home, Windows XP Professional, Windows NT 4 SP4, Windows Server 2003, Windows Vista, or later
- ✔ At least 512MB of total RAM installed on your computer (you get better performance with more)
- ✔ CD/DVD-ROM drive
- ✔ Sound card
- ✔ VGA monitor capable of displaying at least 256 colors, or better
- ✔ Modem with a speed of at least 56 Kbps or greater

If you need more information on the basics, check out other books published by Wiley Publishing, Inc. at www.wiley.com and www.dummies.com.

Using the CD with Microsoft Windows

Well, it's probably high time you got out the nifty CD that came with this book and readied its programs for use. The following sections cover not only installation, but also burning an ISO image to CD as necessary.

Installing the DART CD applications

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, choose 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 to Linux users: You will need to follow the directions for your particular system to mount and view the contents of the CD-ROM.

2. Read through the license agreement, and then click the Accept button if you want to use the CD.

After you click Accept, the License Agreement window won't appear again.

3. When the CD interface appears, you can begin using it to install the programs and run the demos with just a click of a button (or two).



To run some of the programs on the CD, you may have to keep the disc inside your CD-ROM drive. This is a good thing. Otherwise a very large chunk of the program would be installed on your hard drive, taking up hard-drive space and possibly keeping you from installing other software.

How to burn an ISO image to CD

To burn ISO image files, you have to use third-party CD-burning software such as Nero, NTI CD Maker, or any other similar tool. ISO image files should not be directly burned to a CD as normal data. All CD burning software will have an option called “Burn Image” or “Record Image to Disc,” or something similar. Use that option to burn the ISO, so that the CD will become a bootable CD. If you need help with burning ISO images to CD, or with ISO burning software, you can visit this page:

www.ntfs.com/iso-burning.htm

What You'll Find on the DART CD

The following sections are arranged by category; they provide a summary of the software and other goodies you'll find on the *Rootkits For Dummies* DART CD. If you need help with installing the items provided on the CD, refer to the installation instructions in the preceding section.



Shareware programs are fully functional, free, trial versions of copyrighted programs. If you like particular programs, register with their authors for a nominal fee — and receive licenses, enhanced versions, and technical support. *Freeware programs* are free, copyrighted games, applications, and utilities. You can copy them to as many PCs as you like — for free — but they offer no technical support. *GNU software* is governed by its own license, included inside the folder of the GNU software. There are no restrictions on distribution of GNU software. See the GNU license at the root of the CD for more details.



Trial, demo, or evaluation versions of software are usually limited either by time or functionality (such as not letting you save a project after you create it). Be sure you know what their restrictions are before you start exploring their capabilities; it may save you some frustration.

Bonus Chapters

The CD includes two bonus chapters that help familiarize you with specific programs you can use to protect your system from rootkits (and other malware). Bonus Chapter 1 focuses on anti-malware utilities and scanners as well as programs that help you detect and remove rootkits. Bonus Chapter 2 is all about utilities for backing up, cleaning the registry, protecting your passwords, and keeping your hard drive clean.

Anti-malware utilities and scanners

Here are some candidates for your anti-rootkit arsenal that variously fight and (ahem) root out invaders — complete with brief descriptions and information on where to find them.

Agnitum Outpost Firewall

Agnitum Outpost Firewall is an easy-to-use, powerful software firewall. You can upgrade the 30-day trial version to permanent if you purchase the software. (See Bonus Chapter 1 for more about this software and installation instructions.)

Version on DART CD: Trialware with a 30-day trial period.

System requirements: Windows 98/ME /2000/XP/2003 Server. TCP/IP network, Pentium (200 MHz or faster), 32MB of RAM or more, 10MB of available hard-drive space.

Sunbelt Kerio Personal Firewall

Sunbelt Kerio Personal Firewall is an award-winning and affordable form of protection with a limited freeware edition.

Version on DART CD: Trialware with a 30-day trial period.

System Requirements: Windows 2000, XP.

Spybot-Search&Destroy

Spybot-S&D is a free — and reliable — anti-spyware program, including a real-time protection feature that monitors the system's key areas (such as the Registry) for changes, alerting you if it finds anything suspicious. This feature is called the Tea-Timer.

Version on DART CD: Freeware.

System Requirements: Microsoft Windows 98/ME/NT/2000/XP/2003/Preinstallation Environment (PE). For more about Microsoft Windows PE, visit this Web site:

www.microsoft.com/licensing/sa/benefits/winpe.msp

LinkScanner Pro 2.0

LinkScanner Pro provides real-time analysis of http traffic streams, website content and search results to protect against a wide range of online threats, including malicious content, phishing, social engineering, and targeted software exploits.

Version on DART CD: Trialware with a 30-day trial period.

System Requirements: Pentium (1.2 GHz or faster), 256MB RAM, Windows 2000 or XP (Home or Pro); a broadband Internet connection is recommended but the software will work with dial-up modems.

Backup and imaging applications

These programs provide backups of data and full-featured copies of your operating system to limit the damage done by a malware attack.

Acronis True Image

Acronis True Image Home is an easy-to-use backup solution; it's for backing up selected files and folders, or entire disks and partitions. The backups can be stored on hard drives, CDs, Flash memory sticks, or other high-capacity media. (See Bonus Chapter 2 for more about this software.)

Version on DART CD: Trialware with a 15-day trial period.

System Requirements: Windows 98/ME/NT/2000/XP. 256MB RAM.

Karen's Replicator

Karen's Replicator is a free, easy-to-use program for making backups at scheduled intervals, automatically. It's free for personal, noncommercial use. To obtain a commercial license — or if you want to use it at work — you'll need a paid license. (See Bonus Chapter 2 for more about this software and installation instructions.)

Version on DART CD: Freeware for personal and noncommercial use.

System Requirements: All Karenware Power Tools, including the Replicator, work on all Windows versions, provided they have the VB6 Runtime files installed.

Sandboxie

Sandboxie isolates your system from the threats of the Internet. (See Bonus Chapter 2 for more about this software and installation instructions.)

Version on DART CD: Freeware.

System Requirements: Sandboxie can be installed on Microsoft Windows XP, Windows 2003 Server and Windows 2000. The 64-bit version of Sandboxie can be installed on Microsoft Windows XP Professional x64 Edition and Windows 2003 Server x64 Edition.

System-analysis programs

These programs are the sleuths of the DART CD; they help you track down clues left by malware and rootkits that may be hiding on your system.

Autoruns

Autoruns can show the contents of different auto-starting locations in the system, to help you hunt for rootkits and malware. (See Bonus Chapter 1 for more about this software and installation instructions.)

Version on DART CD: Freeware.

System Requirements: Autoruns works on all Windows versions including the 64-bit versions of XP and 2003.

Process Explorer

Process Explorer is a process viewer and manager that can examine processes and show you how they're using CPU resources, memory, DLLs, and modules. (See Bonus Chapter 1 for more about this software and installation instructions.)

Version on DART CD: Freeware.

System Requirements: Process Explorer works on all Windows versions, including 64-bit versions and Vista.

Rootkit-detection-and-removal applications

These programs are what you need for getting tough on rootkits: They help you find the invaders and get rid of them.

DarkSpy

DarkSpy is a very powerful rootkit-detection-and-removal tool, using multiple means of detection — its own, as well as the popular techniques of other applications. It scans processes, drivers, files, ports, and the Registry. DarkSpy excels at detecting kernel-mode rootkits.

Version on DART CD: Freeware.

System Requirements: Windows NT, 2000, and XP.

GMEER

GMEER can find many of the most difficult rootkits by scanning processes, modules, Windows services, and more. GMEER is updated to keep pace with new rootkits.

Version on DART CD: Freeware.

System Requirements: Windows NT, 2000, and XP.

Rootkit Revealer

Rootkit Revealer finds hidden or rootkit files by checking for discrepancies in your system. (**Note:** Rootkit Revealer does not *remove* malware files, and your scan results may contain entries that are not rootkit-related. See Bonus Chapter 1 for more about this software and installation instructions.)

Version on DART CD: Freeware.

System Requirements: Windows NT 4 and higher.

UnHackMe

UnHackMe can detect and completely remove rootkits automatically. Its real-time monitor can prevent the installation of rootkits. (See Bonus Chapter 1 for more about this software and installation instructions.)

Version on DART CD: Trialware with a 30-day trial period.

System Requirements: Windows NT4/2000/XP (compatible with Windows XP SP2)

Password protectors and generators

Here's some help with creating tougher passwords and keeping them uncracked — basic to enhanced security.

Advanced Password Generator

The included version can generate two random and secure passwords at a time, each one four characters long. (See Bonus Chapter 2 for more about this software and installation instructions.)

Version on DART CD: Shareware.

System Requirements: Pentium, Windows 9x/Me/NT/2000/XP/2003, 16MB RAM minimum.

Any Password

Use Any Password to store passwords, user IDs, and related information in encrypted database files protected by a master password. (See Bonus Chapter 2 for more about this software and installation instructions.)

Version on DART CD: Freeware edition.

System Requirements: Any Password works on all Windows versions.

Downloading tools for compromised hard drives

If you're in the midst of trying to recover from a rootkit attack, take heart: Chapters 14 and 15 provide profiles of some major software tools, including some in Bonus Chapter 2 for erasing and repairing compromised hard drives — Eraser, KillDisk, and SpinRite — as well as links to the Web sites where you can get them.

Troubleshooting

We've tried our best to compile programs that work on most computers with the minimum system requirements. Alas, just as "your mileage may vary," your computer may differ; 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 the installation or running of a program. If you get an error message, 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 costly step. However, adding more memory can really help the speed of your computer and allow more programs to run at the same time. This may include closing the CD interface and running a product's installation program from Windows Explorer.
- ✔ **Call Customer Care:** 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 techsupport.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.

Index

• Numerics •

- 16-bit architectures, 105
- 32-bit architectures, 105
- 64-bit architectures
 - GMER working with, 289
 - understanding, 105

• A •

- Abrams, Lawrence (Bleeping Computer Web site owner), 349
- access
 - backdoors giving, 12–13
 - limiting/controlling physical, 140
 - need to limit, 139–140
- account logon events, auditing, 124–125
- account management, auditing, 125
- AccuHash 2.0, verifying system file integrity with, 243
- Acronis True Image
 - on DART CD, 359
 - using, BC14
- Active@ Kill Disk, using, BC22–BC23
- ActiveX
 - blocking, 73–76
 - controls in Internet Explorer 7, 78
 - danger of accepting, 65
- Ad-Aware SE Personal, using, BC2–BC3
- administrative access
 - to disable System Restore, 325
 - rootkits needing, 153
- ADS (alternate data stream), 344
- Advanced Password Generator
 - on DART CD, 362
 - using, BC18–BC19
- Agnitum Outpost Firewall
 - on DART CD, 358
 - using, BC3–BC4
- AIM (AOL Instant Messaging) network, worm infecting, 69
- Answers that work process database, 199
- AntiHookExec
 - benefits of using with other tools, 263–264
 - installing, 263
 - overview, 262
 - user-friendliness of, 245
 - using, BC8
 - using Autoruns with, 265–268
 - using Process Explorer with, 268–269
 - using HijackThis with, 264–265
- anti-malware Real Time Monitoring, instant messaging with, 68
- anti-malware utilities. *See also specific utilities*
 - on DART CD, 358–359
 - recommended, BC2–BC7
- anti-spyware software, 97. *See also specific software*
- anti-trojan software, 98–99. *See also specific software*
- antivirus software, 98
- Any Password
 - on DART CD, 362
 - using, BC19–BC20
- AOL Instant Messaging (AIM) network, worm infecting, 69
- API HookCheck (SIG^2), 164
- API hooking, overview, 159
- Applnit_DLLs injection, functioning of, 165–166
- Application log (Windows), function of, 206
- application programming interface (API)
 - kernel-mode rootkits using, 156
 - user-mode rootkits using, 155
- application-based firewalls, 91. *See also software firewalls*
- applications. *See software*
- Apropos rootkits, 239, 340–341
- archiving, event logs automatically, 208–210
- attachments, guidelines for safe, 66–68

- auditing
 - categories of securities events available for, 124–126
 - overview, 119–120
 - resources for, 120
 - turning on event logging, 121–122
 - turning on security, 122–124
- auditing policies
 - configuring and enabling, 28
 - instant messaging with, 68
- Aumha Web site, 348
- AutoComplete (Internet Explorer), disabling, 77
- automatic updates, setting up, 105–106
- AutoPlay (Windows), disabling on external drives/devices, 59
- AutoRun (Windows)
 - disabling, 58–59
 - working without, 60
- Autoruns (Sysinternals)
 - on DART CD, 360
 - detecting persistent rootkits with, 246–247
 - editing startup list with, 49–50
 - using, BC8–BC9
 - using AntiHookExec with, 265–268
- AVG Anti-Spyware Free, using, BC4

• B •

- backdoor keyloggers, 289–290. *See also* keyloggers
- backdoors
 - allowing access through ports, 183
 - overview, 12–13
 - rootkits opening, 23
- backing up
 - importance of maintaining, 29
 - preparing recovery discs for, 147–148
 - recommended software for, BC14–BC16
 - Registry, 42–43
 - software available for, 304
 - software on DART CD for, 359–360
 - storing after, 304
 - with Windows Backup Utility, 44–46
- bad sectors
 - preparing, 327
 - rootkits hiding in, 326

- BadRKDemo rootkit, DarkSpy removing, 280–282
- bandwidth, measuring use of, 223–224
- BartPE (Bart Preinstallation Environment), 236–237
- baseline
 - using, 146–147
 - value of, 146
- BIOS
 - changing boot order in, 329–331
 - rootkits residing in, 324
- BitDefender antivirus software, 238
- blackhat hacker, 337
- BlackLight (F-Secure)
 - running with Rootkit Revealer, 246
 - scanning for rootkits with, 251–253, 340
 - user-friendliness of, 245
- black-market groups, using malware, 18
- Bleeping Computer Startup Programs Database, 199
- Bleeping Computer Web site, 348–349
- boot order, changing in BIOS, 329–331
- bootable CDs
 - Linux, 316
 - Microsoft, 235–236
 - non-Microsoft Windows, 236–238
- bootsector viruses, scanning for, 60
- browsers
 - configuring securely, 29
 - using alternate, 28
- bundling, installing spyware through, 14–15
- business plans, creating physical security, 143–144
- Butler, Jamie (*Rootkits: Subverting the Windows Kernel*), 177

• C •

- CastleCops Security Forum, determining toolbar legitimacy with, 14
- CastleCops Security Professionals, 349–350
- CastleCops Services Lists, 199
- CastleCops StartupList, 199
- CCleaner, using, BC17

- CDs. *See also* Dummies Anti-Rootkit Toolkit (DART)
- burning ISO image to, 321–322, 357
 - Linux boot, 316
 - making computers boot from, 321
 - Microsoft bootable, 235–236
 - non-Microsoft bootable, 236–238
- Cermak, Mike (Tech Support Guy), 353–354
- chat clients, monitoring instant messaging accounts with, 68
- checklist for improving security, 118–119
- checksums, assessing file integrity with, 240
- ChrisRLG (Malware Removal Web site owner), 351
- cleaners, recommended for Registry and system, BC17–BC18
- cloaked rootkits, defined, 232
- computer
- activity, comparing with bandwidth usage, 224
 - needing new, 115
- computer privileges, understanding, 153.
See also user accounts
- connections, disabling network before cleaning rootkits, 319–320
- Content Zones (Internet Explorer), adjusting, 74–76
- contests, avoiding online, 67
- Coyote, Tom (Tom Coyote Security Forum owner), 354
- CPU cycles, monitoring, 228–229
- cross-diff comparison, using different rootkits-detection tools, 234
- **D** •
- DarkSpy
- analyzing Registry with, 279–280
 - comprehensiveness and user-friendliness of, 245
 - on DART CD, 361
 - detecting/removing rootkits with, 277–278
 - evaluating process activity with, 227–228
 - overview, 276–277
 - Registry Analyzer, 278
 - removing difficult rootkits with, 280–282
 - using for port-to-process mapping, 190
- DART. *See* Dummies Anti-Rootkit Toolkit
- data
- controlling flow of with packets, 199–200
 - tracking suspicious flow of, 191
- Data Sentinel, verifying system file integrity with, 244
- databases
- creating security settings, 131–133
 - researching process, 198
- dd, creating hard drive image with, 316–317
- dd for Windows, copying RAM dump with, 313–314
- DDoS (Distributed Denial of Service)
- network of zombies used in, 18
 - object of attacks, 23
- defragmenting, hard drive, 53–57
- dialers, overview, 12
- Diamond CS, Port Explorer, 190–191, 202–205, 293–300
- direct kernel object manipulation (DKOM), 171–173, 341, 342
- Direct Revenue LLC, litigation against, 309
- directory service access, auditing, 125
- Diskeeper Pro (Executive Software), 57
- Distributed Denial of Service (DDoS)
- network of zombies used in, 18
 - object of attacks, 23
- DLL injection
- Applnit_DLLs, 165–166
 - detecting with IceSword, 290–291
 - detecting with Process Explorer, 291–293
 - functioning of, 164
 - overview, 159
- DLLs (dynamic link libraries)
- kernel and user, 161–163
 - overview, 160, 162
 - rootkits targeting, 160
 - user-mode rootkits using, 155
- domain, defined, 125
- domain controller, defined, 125
- double filename extensions, viewing, 12
- downloading
- guidelines for safe, 65
 - using scanners before, BC2
- drive-by downloads, installing spyware by, 15–16

drivers, installing as rootkits, 166–168
 Dummies Anti-Rootkit Toolkit (DART)
 anti-malware utilities and scanners with,
 358–359
 backup and imaging software on, 359–360
 CD contents, 357–358
 installing CD with Microsoft Windows,
 356–357
 password protectors and generators
 on, 364
 rootkit-detection-and-removal software
 on, 361–362
 system requirements for, 355–356
 system-analysis software on, 360–361
 troubleshooting, 363
 dynamic link libraries. *See* DLLs

● E ●

Easter eggs, backdoors installed as, 13
 EAT (Export Address Table), as avenue to
 DLLs, 162
 eEye BootRoot, 343
 Elite toolbar, 339–340
 e-mail
 guidelines for safe, 66–68
 rootkits facilitating spam, 23
 EnCase, forensic assistance from, 318
 Encrypting File System (EFS), safe surfing
 with, 64
 End User License Agreements (EULAs),
 65, 339
 Eraser, using, BC21–BC22
 Eshelman, James A. (Aumha owner), 348
 Ethereal, sniffing hackers with, 205
 EULalyzer, protecting from spyware, 15
 evaluating Web sites safety with, 81–82
 Event Log Explorer, advantages of, 217–219
 event logging. *See also* auditing
 overview, 119–120
 turning on, 121–122
 event logs
 automatically archiving, 208–210
 changing default size, 207–208
 inspecting with Event Log Explorer,
 217–219
 inspecting with Event Viewer, 210–213

 inspecting with MonitorWare, 219–222
 monitoring for rootkits clues, 180
 overview, 207
 types of, 206
 Event Viewer (Windows)
 accessing, 206–207
 evaluating inspection results in, 213–214
 filtering event log data with, 214–216
 finding rootkits with, 126
 inspecting event logs with, 210–213
 upgrading to Event Log Explorer from,
 217–219
 events
 categories available for auditing, 124–126
 filtering by type of, 214–216
 evidence
 collecting, 304
 collecting RAM dump to USB flash drive,
 312–316
 guidelines for preserving, 310–312
 hiring professional to analyze, 317–318
 tracking perpetrators with, 307–308
 executable files, hidden, 11–12
 Executive Software (Diskeeper Pro), 57
 Export Address Table (EAT), as avenue to
 DLLs, 162
 external media, scanning for bootsector
 viruses, 60

● F ●

false security alerts, encouraging purchase
 of malware programs, 17
 FanBot, 343–344
 Farmer's Boot CD (FBCD), 316
 file analysis services, availability of, 305
 FileAlyzer, verifying system files with,
 240–243
 file-integrity checks, recommending, 145
 Filemon (Sysinternals)
 tracking forensic tool changes with, 312
 tracking outbound access with, 197
 filename extensions, viewing, 12
 files
 backing up with Windows Backup Utility,
 44–46
 checking for legitimacy of, 257

- scanning before opening, 67–68
- tracking forensic tool changes to, 312
- filtering, event log data, 214–216
- FIRE (Forensic and Incident Response Environment) Bootable CD, 316
- firewall logs
 - examining for Internet access attempts, 193–194
 - identifying process ID associated with identified part, 195
 - identifying process with, 195
 - identifying processes loaded by `svchost.exe` with PID 984, 196–197
 - monitoring for rootkits clues, 180–181
- firewalls
 - on DART CD, 358
 - functioning of, 83–84
 - hardware, 84–90
 - importance of having, 82
 - importance of using, 28
 - improving, 118
 - preventing drive-by downloads, 15–16
 - software, 90–93
 - understanding, 83
 - Windows XP, 93–95
- firmware, understanding, 324
- Forensic and Incident Response Environment (FIRE) Bootable CD, 316
- Forensics Acquisition Utilities, 314
- format and reinstall
 - changing boot order in BIOS, 329–331
 - for Microsoft Windows XP, 331–332
 - overview, 327–328
 - preparing for, 328–329
 - rootkits evading, 325–327
 - running rootkit detection software after, 333
 - weighing option of, 305–307
- forums
 - asking help from, 234
 - CastleCops Security Forum, 14
 - helping with Rootkit Revealer, 250–251
 - Malware Complaints, 309
 - on security Web sites, 348–354
- freeware, versus shareware, 357–358.
 - See also specific freeware*

- F-Secure, BlackLight
 - running with Rootkit Revealer, 246
 - scanning for rootkits with, 251–253, 340
 - user-friendliness of, 245
- FU rootkit, 341–342
- FUto rootkit, 175–176, 342

● G ●

- Geeks to Go Web site, 350–351
- Genie-Soft Backup Manager Home, using, BC15
- Gibson Research Corporation
 - listing ports used by trojans, 186
 - SpinRite, 327
- Gladiator Security Forum, 351
- GMER
 - comprehensiveness and user-friendliness of, 245
 - on DART CD, 361
 - detecting/removing rootkits with, 284–286
 - enabling system monitoring/tracing in, 286–287
 - evaluating process activity with, 227–228
 - overview, 283
 - Registry feature, 288
 - using in Safe mode, 287–288
 - working with 64-bit architectures, 289
- gray-market groups, using malware, 18
- Greatis
 - Application Database, 199
 - UnHackMe, 245, 260–261, 361–362, BC12
- Group Policy Objects (GPO)
 - applying to networks, 139
 - importing security templates into, 138–139

● H ●

- HackerDefender rootkit, 185, 338
- hackers
 - catching with sniffers, 200–201
 - lack of intelligence of, 26
 - tracking, 307–308
 - using sniffers, 200
 - ways of using ports, 185–186
- HackerWatch (McAfee, Inc.), 182

- hard drives
 - choosing to reformat, 305–307
 - cleaning with Windows Disk Cleanup Utility, 51–52
 - copying infected, 304
 - creating image of infected, 315–316
 - defragmenting, 53–57
 - downloads available for
 - compromised, 364
 - firmware for, 324
 - malware using space on, 222–223
 - partitioning and formatting, 331–332
 - preparing for reformatting of, 328–329
 - problems of junked up, 46
 - rootkits evading reformatting of, 325–327
 - hard-drive erase and repair utilities, recommended, BC21–BC24
 - hardware
 - needing new computers, 115
 - obtaining updates for, 110–112
 - hardware firewalls
 - Network Address Translation (NAT) on, 87
 - overview, 84–86
 - port blocking/port stealthing with, 87–89
 - Stateful Packet Inspection component of, 89–90
 - hash values, assessing file integrity with, 240
 - Hayes, Bert (*Snort For Dummies*), 182
 - Healan, Mike (SpywareInfo Web site owner), 353
 - Helix Bootable CD, 316
 - heuristics
 - on scanners, 100, 244
 - VICE using, 270
 - HijackThis (HJT), using AntiHookExec with, 264–265
 - HIPS (Host-Intrusion Prevention Software)
 - GMER as, 286–287
 - monitoring for rootkits clues, 183
 - Hogland, Greg (*Rootkits: Subverting the Windows Kernel*), 177
 - hooking
 - direct kernel object manipulation as
 - alternative to, 171–173
 - EAT, 162–163
 - function of, 158–159
 - IAT, 161–162
 - inline, 163–164
 - Interrupt Descriptor Table, 171
 - kernel inline, 170
 - overview, 157–158
 - percentage of threats employing, 245
 - SYSENTER, 170, 344
 - System Service Descriptor Table, 168–169
 - virtual memory manager, 176
 - hooks
 - installing drivers as rootkits, 166–168
 - privileged, 166
 - types of, 159
 - Host-Intrusion Prevention Software (HIPS)
 - GMER as, 286–287
 - monitoring for rootkits clues, 183
 - HOSTS file
 - applying updated, 29
 - using, 72–73
 - Howes, Eric (SpywareWarrior Web site owner), 353

● I ●

 - IANA (Internet Assigned Number Authority), coordinating port assignments, 184–185
 - IAT (Import Address Table), as avenue to DLLs, 161–162
 - IceSword
 - comprehensive nature of, 245
 - destroying rootkits with, 257–258
 - detecting keyloggers with, 290–291
 - detecting rootkits changes with, 253–255
 - evaluating process activity with, 227
 - functions of, 259–260
 - interpreting scan results of, 256–257
 - overview, 253
 - using, BC9–BC10
 - using for port-to-process mapping, 189–190
 - IDS (Intrusion Detection Systems)
 - importance of implementing, 29–30
 - logs, monitoring for rootkits clues, 182
 - IDT (Interrupt Descriptor Table), hooking, 171

- illegal material, rootkits hiding, 23
 - IM (instant messaging), guidelines for
 - safe, 68
 - imaging software
 - on DART CD, 359–360
 - for infected hard drives, 304
 - recommended, BC14–BC16
 - Import Address Table (IAT), as avenue to
 - DLLs, 161–162
 - infection
 - guidelines for preserving evidence of, 310–312
 - options for dealing with, 303–305
 - reformatting and reinstalling “cure,” 305–307
 - tracking perpetrator of, 307–308
 - inline hooking, 163–164, 170
 - instant messaging (IM), guidelines for
 - safe, 68
 - Intermix Media, litigation against, 309
 - Internet Assigned Number Authority (IANA), coordinating port
 - assignments, 184–185
 - Internet browsers
 - configuring securely, 29
 - using alternate, 28
 - Internet connection
 - danger of leaving on, 64
 - defragmenting with, 56
 - Interrupt Descriptor Table (IDT),
 - hooking, 171
 - Internet, disconnecting infected machines
 - from, 324
 - Interrupt service routines (ISRs), rootkits
 - misdirecting from, 171
 - Intrusion-Detection System (IDS)
 - importance of implementing, 29–30
 - logs, monitoring for rootkits clues, 182
 - Intrusion Prevention Systems (IPS),
 - importance of implementing, 29–30
 - inventory procedures, protecting
 - equipment with, 143
 - ISO image, burning to CDs, 321–322, 357
 - ISRs (Interrupt service routines), rootkits
 - misdirecting from, 171
- J •**
- Java, blocking, 73–76
 - JavaScript, blocking, 73–76
 - Jotti, checking for file legitimacy with, 257
- K •**
- Karen’s Replicator
 - on DART CD, 360
 - using, BC15
 - Kaspersky Antivirus version 5.0, Rootkit Revealer interacting with, 249–250
 - kernel DLLs, rootkits targeting, 161–163
 - kernel inline hooking, 170
 - kernel patching, 169
 - kernel-mode rootkits
 - reformatting and reinstalling after, 306
 - summary of, 154
 - types of hooks used, 159
 - versus user-mode rootkits, 155–156
 - kernels, drivers allowing access to, 167–168
 - keyloggers
 - detecting with IceSword, 290–291
 - detecting with Process Explorer, 291–293
 - overview, 289
 - rootkits enabling, 23–24
 - types of, 289–290
 - Kleiman, Dave (Forensic Network Advisor), 318
- L •**
- Laudanski, Paul and Robin (CastleCops Security Professionals Web site owners), 349
 - limited-access user accounts
 - establishing, 70–71
 - safe surfing with, 64
 - using, 27
 - using on networks, 141
 - value of, 141
 - links, embedded within instant-messaging
 - text, 68
 - LinkScanner Pro 2.0, on DART CD, 359

- Linux Knoppix, 316
 - LinuxDefender Live! CD, booting to, 238
 - litigation
 - pursuing, 308
 - against rootkit creators, 308–309
 - Local Security Policy Editor, editing
 - policies and configuring security with, 126–127
 - Log Parser utility, investigating event logs
 - with, 220–222
 - logical bad sectors
 - preparing, 327
 - rootkits hiding in, 326
 - logon attempts, evaluating unsuccessful, 213–214
 - logon events, auditing, 125
 - logs. *See also* event logs; firewall logs
 - monitoring for rootkits clues, 180–183
 - security logs, configuring and enabling, 28
 - sniffer logs, 181–182
- **M** ●
- malicious adware
 - understanding, 13–14
 - ways of installing, 14–16
 - malware
 - chasing dodging, 192
 - ensuring rootkit survival, 152
 - exploiting rootkits, 22
 - finding, 304–305
 - overview, 9–10
 - purpose of, 16–19
 - symptoms of presence of, 31–32
 - types of, 10–16
 - using device drivers, 167
 - Malware Complaints forum, 309
 - Malware Removal Web site, 351
 - master boot record (MBR), 324
 - McAfee, Inc.
 - HackerWatch, 182
 - SiteAdvisor, 81–82, 198
 - MD5summer, verifying system file integrity
 - with, 244
 - Metasploit, maintaining list of native API
 - entries, 169
 - Microsoft Corporation
 - finding updates/patches from, 105
 - getting automatic updates from, 105–106
 - obsolete versus supported systems of, 103
 - offering patches/updates, 103–104
 - responding to rootkit threat, 20
 - Microsoft Inside-the-box GhostBuster, 273
 - Microsoft Internet Explorer
 - adding Web sites to Trusted zone, 76
 - adjusting Content Zones in, 74–76
 - disabling AutoComplete in, 77
 - using version 7, 77–79
 - vulnerability of, 28
 - Microsoft Malicious Software Removal Tool (MSRT), scanning for rootkits with, 261–262, 338
 - Microsoft Management Console (MMC)
 - comparing current security system and template with, 128–130
 - customizing security templates for networks with, 136–137
 - Microsoft Newsgroups, getting help
 - from, 352
 - Microsoft Strider GhostBuster
 - inside-the-box, 273
 - overview, 273
 - WinPE, 274
 - Microsoft TechNet: Events and Errors for the Windows and the Windows Server System Web site, 120
 - Microsoft Telnet
 - disinfecting RATs with, 298–299
 - overview, 297
 - Microsoft Update
 - implementing, 107–109
 - installing, 109–110
 - overview, 106–107
 - preparing for, 107
 - resources on, 112–113
 - Microsoft Virtual PC, 145
 - Microsoft Windows
 - accessing Event Viewer in, 206–207
 - filtering event log data with Event Viewer in, 214–216
 - inspecting event logs with Event Viewer in, 210–213

- installing DART CD with, 356–357
 - removing unused components of, 52–53
 - troubleshooting with, 206
 - Microsoft Windows 2003 Server, turning on
 - security auditing in, 122–124
 - Microsoft Windows Access Control Mechanisms
 - customizing security templates for networks with, 135–139
 - editing policies and configuring security with, 126–127
 - overview, 126
 - testing system against security templates with, 127–135
 - Microsoft Windows Backup Utility, preparing recovery discs with, 147
 - Microsoft Windows Disk Cleanup Utility
 - cleaning hard drive with, 51–52
 - removing unused Windows components/installed programs/system restore points with, 52–53
 - Microsoft Windows Update
 - implementing, 107–109
 - overview, 106–107
 - preparing for, 107
 - resources on, 112–113
 - Microsoft Windows XP
 - firewall, 93–95
 - installing, 331–332
 - Microsoft Windows XP Backup Utility
 - backing up files with, 44–46
 - installing, 44
 - Microsoft Windows XP Pro, turning on
 - security auditing in, 123–124
 - Microsoft WinPE Strider GhostBuster, 274
 - Microsoft WinPE (Windows Preinstallation Environment), booting to, 235–236
 - Microsoft's Security Monitoring and Attack Detection Planning Guide Web site, 120
 - MMC (Microsoft Management Console)
 - comparing current security system and template with, 128–130
 - customizing security templates for networks with, 136–137
 - MonitorWare, investigating event logs with, 219–222
 - Mozilla Firefox
 - advantages of surfing with, 80–81
 - securing, 76–77
 - MSCONFIG, editing startup list with, 48–49
 - MSRT (Microsoft Malicious Software Removal Tool), scanning for rootkits with, 261–262, 338
 - MyFip rootkit, 342–343
- *N* ●
- native API hooking, 169
 - Netbus 1.60
 - disinfecting with Telnet, 298–299
 - disinfecting with Visual Basic program, 299–300
 - tracing with Port Explorer, 293–298
 - Netstat, using for port-to-process mapping, 187–188
 - Network Address Translation (NAT), with hardware routers, 87
 - network-address translation (NAT) capability, routers with, 192
 - networks
 - cleaning of rootkits, 233
 - customizing security templates for, 135–139
 - disabling connections before cleaning rootkits, 319–320
 - monitoring for rootkits, 179–180
 - obtaining updates for, 110–112
 - options for dealing with infections on, 303–304
 - protecting ports of, 183–191
 - rebooting in disinfecting versus reinstalling operating system on, 306
 - using limited-access accounts on, 141–142
 - watching logs for rootkits clues, 180–183
 - New York State, litigation against adware companies by, 309
 - NIDS (Network Intrusion-Detection System) log, monitoring for rootkits clues, 182
 - NOD32 Antivirus, using, BC5
 - Norton Ghost, creating hard drive image with, 315–316
 - NTFSShider, 339
 - NTI Backup NOW!, using, BC15

• 0 •

- object access, auditing, 125
- online contacts, choosing carefully, 62
- Opera, securing, 76–77
- operating systems (OS)
 - choosing to reinstall, 305–307
 - importance of securing, 30
 - obsolete, unsupported, 103
 - preparing for reinstallation of, 328–329
 - rootkits evading reinstallation of, 325–327
 - rootkits hiding in, 151–153
 - rootkits infecting, 234–235

• p •

- packets, controlling data flow with, 199–200
- packet-sniffers, rootkits enabling, 23–24
- page fault, 346
- page-fault handler, 346
- parental control programs, versus spyware, 14
- Password Safe, using, BC20–BC21
- passwords
 - creating and storing, 29, 69–70
 - creating free e-mail accounts, 66
 - hackers deciphering weak, 17
 - recommend protectors and generators for, BC18–BC21
- patches
 - automatic Microsoft, 105–106
 - finding Microsoft, 105
 - Microsoft offering, 103–104
 - reasons for, 104
 - staying current with, 29
 - understanding, 102
- patching, miscellaneous software, 113–115
- peer-to-peer (P2P)
 - file sharing programs, spyware attached to, 15
 - networks, downloading from, 65
- persistent rootkits, detecting, 246–247
- personal information
 - hackers obtaining, 17
 - keyloggers collecting, 24

- pe386, 344
- phish, 343
- physical access
 - importance of limiting/controlling, 142
 - limiting/controlling, 140
- physical security, creating plan for, 143–144
- pings, sending out, 88
- plans, creating physical security, 143–144
- platforms, rootkits specific to, 21
- Pocket KillBox, using, BC17–BC18
- policy changes, auditing, 125
- pornography
 - avoiding, 64
 - drive-by downloads from, 15
- port blocking, with hardware routers, 87–89
- Port Explorer (Diamond CS)
 - sniffing capabilities of, 202–205
 - tracing RATs with, 293–300
 - using for port-to-process mapping, 190–191
- port stealthing, with hardware routers, 87–89
- ports
 - checking with port-to-process mapping, 186–191
 - identifying process ID associated with identified, 195
 - identifying process through, 195
 - overview, 183–185
 - ways hackers use, 185–186
- port-to-process mapping
 - checking ports with, 186–187
 - tracking suspicious data flow, 191
 - using DarkSpy for, 190
 - using Netstat for, 187–188
 - using Port Explorer for, 190–191
 - using TCPView, 188–189, 189–190
- preview panes, disabling e-mail, 66
- Privacy Policies, guidelines for safe downloading with, 65
- process activity
 - evaluating with DarkSpy and GMER, 227–228
 - evaluating with IceSword, 227
 - evaluating with Process Explorer, 226–227
 - evaluating with Task Manager, 224–226

- Process Explorer
 - on DART CD, 360–361
 - detecting keyloggers with, 291–293
 - evaluating process activity with, 226–227
 - identifying processes loaded by
 - `svchost.exe` with PID 984, 196–197
 - using, BC10
 - using AntiHookExec with, 268–269
 - The Process Library, 199
 - process tracking, auditing, 125
 - processes
 - identifying, 195–197
 - mapping recipient IP address with
 - reverse DNS search, 198
 - researching databases of, 198
 - professionals, hiring to analyze evidence, 317–318
 - programs. *See* software
 - PspCidTable, FUTO rootkits altering, 175–176
 - P2P (peer-to-peer)
 - file sharing programs, spyware attached to, 15
 - networks, downloading from, 65
 - puppet masters, rootkits delivering, 22–23
- R •**
- RAIDE, scanning for rootkits with, 275–276, 345
 - random-access memory (RAM), dumping to USB flash drive, 312–314
 - RATs (remote-access trojans)
 - accessing through ports, 183
 - commandeering computers, 17–18
 - disinfecting with Telnet, 298–299
 - reinstalling operating system for, 306–307
 - tracing with Port Explorer, 293–298
 - using ports, 185
 - writing Visual Basic program to disinfect, 299–300
 - reboots, planning after system compromise, 320–322
 - recognition
 - overview, 30–31
 - of problems not from malware, 33
 - signs of malware, 31–32
 - recovery, planning for, 33–34
 - recovery discs, preparing, 147–148
 - Red Screen of Death (RSOD), 343
 - Registrar Registry Manager, using, BC18
 - Registry
 - analyzing with DarkSpy, 279–280
 - backing up, 42–43
 - cleaning, 57–58
 - editing with GMER, 288
 - recommended cleaners for, BC17–BC18
 - tracking forensic tool changes to, 313
 - Registry Editor, backing up Registry with, 42–43
 - Regmon (Sysinternals), tracking forensic tool changes with, 313
 - remote transfers, installing spyware by, 15
 - remote-access trojans (RATs)
 - accessing through ports, 183
 - commandeering computers, 17–18
 - disinfecting with Telnet, 298–299
 - reinstalling operating system for, 306–307
 - tracing with Port Explorer, 293–298
 - using ports, 185
 - writing Visual Basic program to disinfect, 299–300
 - resistance
 - becoming intelligent computer user for, 26–27
 - need for, 26
 - security measures recommended for, 27–30
 - resources
 - on developing kernel device drivers, 168
 - on hard drive operation, 326
 - lists of ports used by trojans, 186
 - for security devices, 143
 - for security software, 144
 - security Web sites, 348–354
 - for software updates/patches, 114
 - for virtual machine software, 145
 - restore point. *See also* System Restore
 - creating, 39
 - removing installed, 52–53
 - restoring from, 41
 - restoring
 - Registry, 43
 - system, 41

- Restricted Groups policy, implementing, 141–142
 - Rootkit Revealer Forum, 250
 - Rootkit Revealer (RKR) (Sysinternals)
 - on DART CD, 361
 - help resources for, 250–251
 - interpreting scan results of, 249–250
 - overview, 247–248
 - running with BlackLight, 246
 - understanding operation of, 248–249
 - user-friendliness of, 245
 - using, BC10–BC11
 - rootkit scanners. *See also specific scanners*
 - function of, 232–233
 - overview, 231–232
 - running after format and reinstall, 333
 - types of, 244–246
 - using from external storage devices, 235–238
 - using in Safe mode, 238–239
 - rootkits
 - cloaked verses on cloaked, 232
 - evading security systems, 21–22
 - history of, 19–20
 - invisible functions of, 151–153
 - kernel-mode versus user-mode, 155–156
 - overview, 19
 - planning defense against, 145–146
 - possible actions against creators of, 308–310
 - reasons for existence, 22–24
 - steps for removing, 233
 - types of, 154
 - Rootkits: Subverting the Windows Kernel* (Hogland, Butler), 177
 - Rootkitty, scanning for rootkits with, 274–275
 - routers, blocking tools with, 192–193
 - RSOD (Red Screen of Death), 343
 - rule-based firewalls, 91. *See also* software firewalls
 - Russonovich, Mark (Sysinternals), 20, 352
 - RxTx, measuring bandwidth use with, 223–224
- S ●
- SA (SiteAdvisor) (McAfee, Inc.), 81–82, 198
 - Safe mode
 - booting into with System Configuration Utility, 229
 - defragmenting in, 55–56
 - rootkits scanning and, 238–239
 - running system restore from, 41–42
 - using GMER in, 287–288
 - Safer-Networking Web site, 243
 - sandbox, defined, BC16
 - Sandboxie
 - on DART CD, 360
 - using, BC16
 - SANS Institute, listing ports used by trojans, 186
 - saving, before reboots, 320
 - scam operations, hackers using systems for, 17–18
 - scanners. *See also* rootkit scanners; *specific scanners*
 - anti-spyware, 97
 - anti-trojan, 98–99
 - antivirus, 98
 - on DART CD, 358–359
 - importance of having, 82
 - importance of using, 28
 - most useful features of, 99–100
 - overview, 95–96
 - type of, 96–97
 - using before Internet downloads, BC2
 - Scott, Charlie (*Snort For Dummies*), 182
 - screen savers, defragmenting with, 56
 - security
 - configuring, 126–127
 - configuring selectively, 134–135
 - creating plan for physical security, 143–144
 - guidelines for improving, 118–119
 - importance of taking physical measures, 142
 - security auditing
 - events, categories available for auditing, 124–126
 - turning on, 122–124

- Security Configuration Manager, editing
 - policies and configuring security with, 126–127
- security devices, protecting equipment with, 143
- security forums
 - asking help from, 234
 - helping with Rootkit Revealer, 250–251
- Security log (Windows), function of, 206
- security logs, configuring and enabling, 28
- security permissions, checking for
 - alteration of, 240
- security policies
 - editing, 126–127
 - setting up, 28
- security settings
 - comparing with security template, 128–130
 - copying from template to template, 137
 - creating database of, 131–133
- security systems, rootkits evading, 21–22
- security templates
 - customizing for networks, 135–139
 - testing system against, 127–135
- Self Monitoring And Reporting Technology (S.M.A.R.T.), helping rootkits hide, 326
- service installations/starts, evaluating for rootkits, 214
- security-analysis utility, making own, 127
- Shadow Walker, 345–346
- shareware, versus freeware, 357–358
- SIG² (API HookCheck), 164
- Sigcheck, verifying system file integrity with, 243
- Silberman, Peter (FUTo developer), 175
- Simovits, listing ports used by trojans, 186
- SiteAdvisor (SA) (McAfee, Inc.), 81–82, 198
- 16-bit architectures, 105
- 64-bit architectures
 - GMER working with, 289
 - understanding, 105
- size, changing event logs default, 207–208
- S.M.A.R.T. (Self Monitoring And Reporting Technology), helping rootkits hide, 326
- sniffer logs, monitoring for rootkits clues, 181–182
- sniffers
 - catching hackers with, 200–201
 - detecting, 201–202
 - Ethereal, 205
 - function of, 24
 - hackers using, 200
 - Port Explorer, 202–205
- Snort For Dummies* (Scott, Wolfe and Hayes), 182
- Snort, monitoring NIDS with, 182
- software. *See also specific software*
 - anti-spyware, 97
 - anti-trojan, 98–99
 - antivirus, 98
 - backup and imaging, BC14–BC16
 - blocking illegal computer use with, 144
 - for detecting/removing rootkits, BC7–BC12
 - editing startup configurations, 48–50
 - false security alerts encouraging
 - purchase of, 17
 - filtering by specific, 214–216
 - hard-drive erase and repair utilities, BC21–BC24
 - recommended anti-malware scanners and utilities, BC2–BC7
 - recommended password protectors and generators, BC18–BC21
 - recommended system and Registry cleaners, BC17–BC18
 - removing installed, 52–53
 - removing unused, 50–51
 - running to many at startup, 47
 - updating/patching miscellaneous, 113–115
 - virtual machine, 144–145
- software firewalls
 - functioning of, 91–92
 - knowing system needs, 95
 - overview, 90–91
 - setting rules for, 92–93
 - Windows XP, 93–95
- SONY Digital Rights Management (DRM)
 - rootkit, discovery of, 20
- spam e-mail, rootkits facilitating, 23
- SPI (Stateful Packet Inspection) component, with hardware firewalls, 89–90

- spies, rootkits acting as, 23–24
- SpinRite (Gibson Research Corporation)
 - described, 327
 - using, BC23–BC24
- Spitzer, Elliott (New York State Attorney General), 309
- spoofed Web sites, drive-by downloads from, 15
- Spybot-Search&Destroy
 - on DART CD, 359
 - inserting Hosts list into HOSTS file, 72–73
 - using, BC6
- spyware
 - understanding, 13–14
 - ways of installing, 14–16
- SpywareInfo Web site, 352–353
- SpywareWarrior Security Web site
 - overview, 353
 - resource on cyber-cor-pirates, 18
- SSDT (System Service Descriptor Table)
 - hooking, 168–169
 - replacing, 170–171
- SSV (System Virginty Verifier), scanning for rootkits with, 270–273
- startup
 - editing list using MSCONFIG, 48–49
 - running too many applications at, 47
- Stateful Packet Inspection (SPI)
 - component, with hardware firewalls, 89–90
- static DLLs, rootkits targeting, 160
- stop errors, evaluating for rootkits
 - activity, 214
- strategies, recommended for rootkit detection and removal, 234
- SubVirt rootkit, 177–178
- Sunbelt Kerio Personal Firewall, on DART CD, 358
- survivable systems, defined, 25
- svchost.exe
 - identifying process started by, 195–196
 - identifying processes loaded by, 196–197
- Symantec, list of top 19 threats from, 245
- symptoms, of malware presence, 31–32
- SYSENTER hooking, 170, 344
- Sysinternals
 - Autoruns, 49–50, 246–247, 265–268, 360, BC8–BC9
 - Filemon, 197, 312
 - PageDefrag, 57
 - Regmon, 313
 - Rootkit Revealer, 245–251, 361, BC10–BC11
 - TCPView, BC11–BC12
- Sysinternals Forum, 352
- system calls, user-mode rootkits using, 155
- system cleaners, recommended, BC17–BC18
- System Configuration Utility, booting into
 - Safe mode with, 229
- system events, auditing, 125
- system files
 - checking for replacement of, 240
 - verifying integrity of, 243–244
 - verifying with FileAlyzer, 240–243
- System log (Windows), function of, 206
- system monitoring, enabling in GMER, 286–287
- system requirements
 - Acronis True Image, 359
 - Advanced Password Generator, 362
 - Agnitum Outpost Firewall, 358
 - Any Password, 362
 - Autoruns, 360
 - BartPE CD, 237
 - checking downloads for, 109
 - DarkSpy, 361
 - DART, 355–356
 - GMER, 361
 - Karen’s Replicator, 360
 - LinkScanner Pro 2.0, 359
 - Process Explorer, 360
 - Rootkit Revealer, 361
 - Sandboxie, 360
 - Spybot-Search&Destroy, 359
 - Sunbelt Kerio Personal Firewall, 358
 - troubleshooting problems with, 363
 - UnHackMe, 361
 - Virtual PC, 145
 - Workstation, 145
- system resources, malware using, 222–223

System Restore
 accessing, 39
 changing drive settings of, 39–40
 clearing out/shutting off, 40–41
 creating restore point, 39
 disabling, 325
 restoring from restore point with, 41
 running from Safe mode, 41–42
 understanding, 38
 working with after system compromise, 323–324

System Service Descriptor Table (SSDT)
 hooking, 168–169
 replacing, 170–171

system service, filtering by, 214–216

system tracing, enabling in GMER, 286–287

System Virginty Verifier (SSV), scanning
 for rootkits with, 270–273

system-analysis software, on DART CD, 360

• T •

Task Manager
 evaluating process activity with, 224–226
 monitoring computer activities with, 223–224
 monitoring CPU cycles with, 228–229

TCPView (Sysinternals), using, BC11–BC12

TCPView, using for port-to-process
 mapping, 188–189

Tech Support Guy Forum, 353–354

32-bit architectures, 105

thread injection, functioning of, 165

Tom Coyote Security Forum, 354

toolbars, adware, 13–14

tools. *See* utilities

training, security awareness, 144

Tripwire, verifying system file integrity
 with, 243

trojanized utilities, replacing system files,
 174–175

trojans. *See also* remote-access trojans
 (RATs)
 effects of, 63
 overview, 11
 using Applnit_DLLs injection, 165
 using ports, 185–186

Trusted zone (Internet Explorer), adding
 Web sites to, 76

Turner, Suzi (SpywareWarrior Web site
 owner), 353

• U •

Ultimate Boot CD For Windows
 (UBCD4Win), booting to, 237–238

uncloaked rootkits, defined, 232

UnHackMe (Greatis)
 on DART CD, 361
 scanning for rootkits with, 260–261
 user-friendliness of, 245
 using, BC12

University of Connecticut (UCONN),
 housing rootkit, 154

unpatched computers, worms effect on,
 11–12

unprivileged hooks, overview, 159

updates. *See also* Microsoft Update;
 Microsoft Windows Update
 automatic Microsoft, 105–106
 finding Microsoft, 105
 importance of, 103
 Microsoft offering, 103–104
 miscellaneous software, 113–115
 for networks and hardware, 110–112
 reasons for, 104
 staying current with, 29
 understanding, 102

USB flash drive, dumping RAM to, 312–314

user accounts
 establishing limited-access, 70–71
 safe surfing with limited-access, 64
 using limited-access, 27
 using limited-access on networks, 141
 value of limited-access, 141

user DLLs, rootkits targeting, 161–163

user-friendliness, of rootkits actors, 245

user-level rootkits, types of hooks
 used, 159

user-mode rootkits
 versus kernel-mode rootkits, 155–156
 summary of, 154

users, filtering event by, 214–216

utilities. *See also specific utilities*
 anti-malware on DART CD, 358–359
 to detect sniffers, 202–205
 hard-drive erase and repair, BC21–BC24
 recommended anti-malware, BC2–BC7
 trojanized replacing system files, 174–175
 for verifying system file integrity, 240–244

• U •

VICE (Virtual Intruder Capture Engine),
 scanning for rootkits with, 269–270
 video-card EEPROM, rootkits residing
 in, 324
 virtual machine (VM) software, 144–145
 virtual memory manager (VMM), hooking,
 176, 345
 virtual operating systems, fooling rootkits
 with, 144–145
 virtual-machine-based rootkits (VMBR),
 function of, 177–178
 Virus Total, checking for file legitimacy
 with, 257
 viruses
 finding, 304–305
 overview, 11
 Visual Basic programs, writing to disinfect
 Netbus, 299–300
 VM (virtual machine) software, 144–145
 VMware Workstation, 145
 vulnerability, to malware, 10

• W •

Web sites
 adding to Trusted zone, 76
 for bootable CDs, 236–238
 for checking for legitimacy of files, 257
 drive-by downloads from spoofed, 15
 evaluating safety of, 81–82
 on hard drive operation, 326

for Linux boot CDs, 316
 for maintaining list of native API
 entries, 169
 for process databases, 198
 resources for phishing, 343
 resources for security auditing, 120
 resources for security devices, 143
 resources for security software, 144
 resources for software updates/
 patches, 114
 resources for virtual machine
 software, 145
 resources listing ports used by
 trojans, 186
 resources on developing kernel device
 drivers, 168
 Safer-Networking, 243
 security, 338, 348–354
 Web surfing
 advantages of using Firefox for, 80–81
 guidelines for safe, 63–64
 importance of practicing safe, 30
 malware controlling, 16–17
 malware tracking, 16
 rootkits tracking, 24
 whoami .exe tool, tracking user access
 with, 141
 Windows File Protection (WFP), protecting
 DLLs, 160
 WinTasks Process Library, 199
 Win32/Hackdoor.B rootkit, using ports, 185
 WMF (Windows Metafile) exploit, installing
 spyware by, 15–16
 Wolfe, Paul (*Snort For Dummies*), 182
 worms, overview, 11

• Z •

zombies, for DDoS, 18
 ZoneAlarm Pro Firewall, using, BC6–BC7

Bonus Chapter 1

Ten (Plus Three) Malware Utilities and Scanners

In This Chapter

- ▶ Ad-Aware SE
 - ▶ Agnitum Outpost Firewall
 - ▶ Ewido Anti-Malware
 - ▶ NOD32 Antivirus
 - ▶ Spybot-Search&Destroy
 - ▶ ZoneAlarm Pro Firewall
 - ▶ AntiHookExec
 - ▶ Autoruns
 - ▶ IceSword
 - ▶ Process Explorer
 - ▶ Rootkit Revealer
 - ▶ Sysinternals TCPView
 - ▶ UnHackMe
-

Elsewhere in *Rootkits For Dummies* we have compared the Internet to a virtual (and very real) battleground, referring to the continuing war with malware and rootkits. The software discussed in this chapter and Bonus Chapter 2 are your weapons and equipment, the tools you need to protect your computer and remove malware and rootkits that try to invade. Your firewall is defensive equipment; most scanners actively detect and remove the presence of enemy infiltration. Most of these powerful programs are discussed in these two bonus chapters.



Some of these applications are included on the *Rootkits For Dummies* DART CD (for a complete list of those, check out the Appendix); we've identified them with this icon. For the others — which are readily downloadable off the Web — we include pointers to help you get hold of them.

Using Internet Downloads

If you do not yet use an antivirus scanner, please install one before you proceed to download and install any further applications. We have done our very best to ensure that the contents of the DART-CD are pristine and pure, but you will need the antivirus most for the programs you will be downloading from the Internet. If you already have an installed and up-to-date antivirus scanner, then you're all set. For those who don't, you will have to trust that the antivirus you install is free of malware — you have to start somewhere.

For the applications listed here that are *not* on the DART-CD, wherever possible we have provided direct download addresses. Type each in your browser's address bar and hit Enter. A box should pop up to ask if you want to open or save the file to your hard drive. Click Save and then choose the location or folder you wish to place it in. Wait for the download to complete.

Using Windows Explorer, navigate to the file you downloaded. Right-click it and select the scan button for your antivirus from the context menu. It will then scan that file and let you know if it's clean. To install it, double-click the file and a wizard will open to help you install. Follow the on-screen instructions. This applies to most of the applications listed here that you'll download online. For the rest, we supply more specific instructions or places where you can get them.

Anti-Malware Utilities and Scanners

These firewalls, antivirus, and anti-spyware applications are essential to preventing rootkits from invading your computer. They all stop the malware that rootkits need in order to gain a foothold on your system. (See Chapter 4 for more information about anti-malware scanners and utilities.)

Ad-Aware SE Personal

Ad-Aware is a popular anti-spyware program that can detect various adware, spyware, and other sorts of malicious programs.

Web site: www.lavasoft.de/software/adaware

Download from: www.lavasoft.de/products/ad-aware_se_personal.php

Version information: Shareware; free for personal and noncommercial use.

Clicking the *Download* button (top of page) at the above address takes you to CNET, one of Lavasoft's download partners. Download the file to the folder or location of your choice. Install it following the instructions in the "Using Internet Downloads" section earlier in the chapter. While installing, you can change the default installation location to one of your choice. After the installation, Ad-Aware can be launched using the icon on the Desktop or through the shortcut in the Start Menu.

To start the scan, click the Scan now button in the left pane. This opens the Scan Options window. Here you can select the type of scan to perform. Select the desired scan type (for example, Full system scan) and click Next. When the scan completes, click Next to go to the Scanning Results window. Here, select the Critical Objects tab and select all the objects that have been detected by Ad-Aware. Finally, click Next to remove the detected infections.



Agnitum Outpost Firewall

Your DART CD includes a trial version of Agnitum Outpost Firewall, an easy-to-use, powerful software firewall.

Web site: www.agnitum.com

Version information: Pro and Free (trialware) versions available. The Free version has a 30-day trial period and is available at www.agnitum.com/products/outpost/download.php.

The Web site provides a forum board where you can get help and ask questions: www.outpostfirewall.com/forum.

To install Agnitum Outpost Firewall, double-click the name of the program on the CD to start the installation. Then follow the on-screen instructions provided by the installation wizard. If you want to change the default location to which you're installing the program, you can do that while installing it.

When installed, Outpost Firewall can be accessed through its System Tray icon, which looks like a question mark (?). Right-clicking the Outpost Firewall icon in the System Tray brings up a menu that provides various options. To open the Outpost Firewall's main window, click the Show button from the menu. In the main window of Outpost Firewall, different sections can be accessed from the links present in the left pane. Selecting an option from the

left pane will result in the display of corresponding information in the right pane. For example, if the Network Activity option in the left pane is selected, Outpost Firewall will display all the programs and processes accessing the network. Additional features of Outpost Firewall, such as “Ad-blocker,” and “Antispyware,” can be accessed by selecting the Plug-ins option in the left pane.

Outpost Firewall has advanced features such as rule-based filtering, application filtering, and more. You can learn about all the features of Outpost Firewall from the PDF user guide, which is available at this Web page:

www.agnitum.com/support/docs.php

AVG Anti-Spyware Free

AVG Anti-Spyware Free is one of the most popular anti-malware applications.

Web site: <http://free.grisoft.com/doc/5390/lng/us/tp1/v5>

Direct Download: <http://free.grisoft.com/softw/70free/setup/avgas-setup-7.5.0.50.exe>

User Manual: http://free.grisoft.com/filedir/doc/AVG_Anti-spyware/User_manual/avg_asw_uma_en_75_4.pdf

Version information: Trialware (full-featured) with a 30-day trial period. When the trial period expires, the program turns into a limited free version.

Double-click the direct download file to start the installation of AVG Anti-Spyware Free. Follow the on-screen instructions to proceed with the installation. You can start the program by double-clicking its icon (present either on the Desktop or Start Menu). Next, to update the definitions, click the Update button in the left pane of the main screen. Then click Start Update to start the update process.

To scan the system, click the Scanner button in the left pane of the main screen, and then click Complete system scan. If it finds any infections, it will pop up a notification. Select Clean, check the boxes beside *Perform action with all infections* and *Create encrypted backup*, and then click OK. When the scan finishes, click Save Report to save the scan log (which can be reviewed later if needed).

NOD32 Antivirus

NOD32 is a critically acclaimed antivirus application. It's very light on system resources and has excellent detection capabilities.

Web site: www.eset.com/products/index.php

NOD32 for Windows: www.eset.com/products/windows.php

Indirect Download: www.eset.com/download/index.php

Choose the circumstance which applies to you from Option One or from Option Two to download the 30-day trial version of NOD32 Antivirus.

Version information: Trialware with a 30-day trial period.

Double-click the installation file you downloaded to start the installation. The installer will prompt you to choose the location to extract the installation files. Click the Extract button to extract the files to the default location. It asks you to choose the type of installation. The options provided are Typical, Advanced, and Expert. It is highly recommended to choose the Typical option. On the later stages of installation, the installer will ask you to configure the automatic update server for NOD32. We recommend the *Choose automatically* option for the update server. If you're using a dial-up Internet connection, put a check mark for the corresponding option in the next screen. After completing the installation, restart the system to save the changes made to it.

NOD32 Control Center can be invoked by simply clicking its System Tray icon. NOD32 has many components — namely, AMON, DMON, EMON, IMON, and NOD32. Each component has a specific functionality. To perform a complete system scan, click the NOD32 text in the Control Center. Then click the Run NOD32 button displayed in the right-side window. After a few seconds, the scanning window will open; select the required scanning targets that are to be scanned by NOD32. After this, click the Scan & Clean button to start the scan.

For a detailed explanation of each of the features available in NOD32, you can download the user guide from their Web site:

www.nod32.com/download/manuals.php



Spybot-Search&Destroy

Spybot-S&D is a free — and reliable — anti-spyware program — including a real-time protection feature that monitors the system's key areas (such as the Registry) for changes, alerting you if it finds anything suspicious. This feature is called the Tea-Timer.

Web site: www.safer-networking.org/en/spybotsd/index.html

Version information: Freeware.

Double-click the name of the program on the CD to start the installation wizard. The installer will ask about the default program language. After you choose the desired language, the installer displays the license agreement. Read and accept the license agreement to continue the installation. In the next screen, the installer will display the default path where the program will be installed. If you wish, change the path, then select the desired components to be installed and proceed with the installation.

Spybot-S&D can be started by clicking its shortcut icon (either on Desktop or in the Start Menu). When Spybot-S&D is invoked for the first time, it displays a Wizard. This Wizard helps you to perform tasks like updating the database of Spybot-S&D and creating a backup of the System Registry. It is important to update the database, because Spybot-S&D will not contain a full-fledged database when installed.

To start a system scan, click the topmost button (labeled Spybot-S&D) and then click Check for problems. After the scan, if Spybot-S&D finds any threats, they can be removed by clicking the *Fix selected problems* button, after selecting the threats that are to be removed. A complete tutorial for Spybot-S&D is at www.safer-networking.org/en/tutorial/index.html.

Spybot-S&D contains a well-organized Help File, where you can find a Tutorial to help you to learn how to use the program, especially important for first-time users. Simply click Help, then Tutorial. The Help files are a rich source of information on malware, including a dictionary, FAQ, and Tool List.

ZoneAlarm Pro Firewall

ZoneAlarm is one of the best firewalls out there. It's available in Pro and Free versions. The Pro version has additional features such as anti-spyware, a cache cleaner, support for expert rules, and anti-rootkit measures.

Web site: www.zonelabs.com:80/store/content/home.jsp

Version information: Trialware with a 15-day trial period.



All of the addresses to the ZoneLabs Web site's download pages are extremely long, too long to include here. You will have to choose the program you want and download it yourself. We have provided instructions for installing the ZoneAlarm Firewall Pro only. If you choose their other programs instead, please refer to ZoneLabs' instructions.

Double-click the file you downloaded from ZoneLabs to start the installation. Follow the instructions given out by the installation wizard. You need to provide a name and an e-mail address while installing. If necessary, you can change the default installation folder to one of your choosing. When the installation is finished, you're greeted with the Configuration Wizard, which helps you to set up and customize the firewall. Follow the onscreen instructions. ZoneAlarm's main window can be accessed by double-clicking its System Tray icon. Different sections of the firewall can be accessed by clicking the respective menu buttons in the left pane.

When you first use the ZoneAlarm Pro, you may get lots of alerts from it. But don't worry — most of the time ZoneAlarm is simply informing you that it has blocked some attack. To disable these kinds of alerts, click the Alerts & Logs button in the left pane of main window. Here in the Main tab, select the Off radio button under the Alert Events Shown option box. By doing this, ZoneAlarm will not inform you when it has blocked an inbound attack. Whenever a program tries to connect to network or Internet, ZoneAlarm will intercept it and will pop up a notification. You can either allow or deny the connection to the program that is trying to connect. If you need ZoneAlarm to "remember" your decision about allowing or denying a connection to a specific program, select the *Remember this decision* option.

You can see what programs are connecting to the network by clicking the Program Control button in the left pane of the main window and then selecting the Programs tab. Finally, to turn on the anti-spyware feature of ZoneAlarm, click the Anti-spyware button in the left pane and then click the On radio button. To know about the advanced options available in ZoneAlarm, you can download the PDF user guide from this Web page:

www.zonelabs.com/store/content/support/za/znaIrmMain.jsp

Applications That Detect and Remove Rootkits

Here we list the most powerful rootkit detectors and removers you can get today. We have provided basic instructions for their use, but to see them in their element please refer to Chapter 9 and the index as they are featured throughout this book. The authors have tested and used every one of these applications; it's all good equipment, tried and true.

AntiHookExec

AntiHookExec is a command-line tool, by which any other program can be launched with all API hooks removed. It searches for API hooks and restores them by patching their DLLs. AntiHookExec is a powerful tool for detecting the presence of rootkits and other forms of malware; for a look at how it works, see Chapter 9.

Web site: www.security.org.sg/code/antihookexec.html

Direct Download: www.security.org.sg/code/AntiHookExec.zip

Version information: Freeware.

Copy the .ZIP archive of AntiHookExec to your hard drive. Extract the contents of the .ZIP archive to a dedicated folder, using the Extract All option present in the context menu (Windows XP and later) or using any third-party file compression tool.

Go to Start Menu ⇨ Run and type **cmd** and press Enter to launch the Command Prompt. Here, in the Command Prompt, navigate to the folder where AntiHookExec.exe is extracted, using the `cd` command. When you're in the folder where AntiHookExec.exe is present, then type the command **AntiHookExec.exe <Program Name>** and press Enter. Here <Program Name> is the name of the program, with its complete path, that you want AntiHookExec to launch.



For illustrated instructions on using AntiHookExec go to this address:

<http://swatrant.blogspot.com/2006/02/detecting-rootkits-using-normal-tools.html>

Autoruns



Autoruns is a free program from Sysinternals that can show the contents of different autostarting locations in the system. Autoruns is much more powerful than the MSConfig tool that's available in Windows.

www.sysinternals.com/utilities/autoruns.html

Web page: www.microsoft.com/technet/sysinternals/utilities/Autoruns.msp

Version Information: Freeware. Note that Microsoft now owns Sysinternals but still provides Autoruns as a freeware program.

Copy the zip archive of Autoruns to your hard drive. Extract the contents of the zip archive to a dedicated folder, using the Extract All option present in the context menu (Windows XP and above) or using any third-party file compression tool.

Autoruns has two versions, a GUI (Graphical User Interface)-based program called Autoruns.exe and a command-line program named Autorunsc.exe. Double-click the Autoruns.exe to start the GUI-based Autoruns. As soon as the program is started, it will scan the system and list all the autostart programs. The program shows various tabs for different autostart locations, and the Everything tab lists all the autostart entries together. You can save this list of autostart entries in a log by choosing the File Menu ⇨ Save option. To minimize the size of the log, you can also choose to hide the default Microsoft autostart entries by checking the Options ⇨ Hide Microsoft entries option.

IceSword

IceSword is currently one of the most advanced anti-rootkit tools. And it's available for free. For a look at IceSword in action, see Chapter 9.

Web site: www.xfocus.org (English and Chinese)

For more instructions and downloads: http://castlecops.com/t156595-How_to_Remove_Rootkits_with_IceSword_English.html

Version information: Freeware.

IceSword's executable file comes in a RAR archive named `IceSword_en1.16.rar`. Copy this archive to a desired location in your hard drive. To extract this archive, you need a third-party file-compression tool such as WinRAR. If you don't already have a tool that can handle .RAR file types, download WinRAR or an equivalent before you use this application. (An equivalent freeware RAR tool is available at the CastleCops online address.)

Extract the IceSword archive to a dedicated folder. Next, double-click `IceSword.exe` to start the program. Clicking the Processes tab in the left pane shows all the running processes in the system. If a rooted (hidden) process is present, it will be shown in red color. Similarly, the Win32 Services tab shows all the Windows services. Here also, a rooted service will be shown in red color. IceSword can kill the processes, when the Kill option is selected from the right-click context menu that's available for each process. Similarly, a service can be stopped or paused by selecting the desired option from the context menu that appears when you right-click any of the services. IceSword also has built-in file-system and Registry browsers. They can be reached by clicking the Registry and File tabs, respectively. The file-system explorer and Registry editor are capable of showing rooted files, folders, and Registry entries.

IceSword requires administrator privileges to run, as it uses a driver to operate in kernel mode. Before killing or deleting any files using IceSword, please make sure that the files are indeed malware.

Process Explorer



Your DART CD includes Process Explorer — a free, powerful process viewer and manager. It can show additional information such as CPU usage history and memory usage of a process, DLLs loaded by a process, and other incriminating details.

Web page: www.microsoft.com/technet/sysinternals/utilities/ProcessExplorer.aspx

Version information: Freeware. Note that Microsoft now owns Sysinternals but still provides Process Explorer as a freeware program.

Copy the .ZIP archive to your hard drive. Extract the contents of the .ZIP archive to a dedicated folder using the Extract All option present in the context menu (Windows XP and above) or using any third-party file compression tool.

Double-click the `Procexp.exe` file to start the program. Process Explorer contains two sub-windows.

- ✓ The top window shows all the running processes in the system.
- ✓ The bottom window shows either DLLs loaded by a selected process or the handles of a selected process, depending on the mode of operation of Process Explorer (which can be changed in the Options menu).

Detailed properties of a process can be obtained by double-clicking that process name in the top window. Process Explorer then shows you additional information, such as DLLs, strings, CPU and memory usage, private bytes, and so on.

Rootkit Revealer



Your DART CD includes Rootkit Revealer from Sysinternals — a free program for finding hidden or rootkit files. It checks for discrepancies in your system to find them. (**Note:** Rootkit Revealer will not remove malware files.)

Web page: www.microsoft.com/technet/sysinternals/utilities/RootkitRevealer.msp

Version information: Freeware. Note that Microsoft now owns Sysinternals, but still provides Rootkit Revealer as a freeware program.

Copy the .ZIP archive to a convenient location on your hard drive. If you have Windows XP or later, right-click the .ZIP archive and choose the Extract All option to extract the files to a folder. You can also use any third-party file-compression tools such as WinZip to extract the files in the archive.

Next, double-click the file `Rootkitrevealer.exe` to start the program and click the Scan button to begin the scan. It may take a while to scan the system. When the scanning is in progress, leave the system idle so as to minimize false positives. When the scan completes, it will show the results in a Window. You can also save the results using the File Menu ⇨ Save option. This will save the log in a file named `Rootkitrevealer.txt`.

Rootkit Revealer uses a driver to perform low-level scans, so you must be logged on as a system administrator to use it. Also, Rootkit Revealer will not run in Safe mode. Keep all the other computing activities as minimal as possible while the Rootkit Revealer is scanning. Disable unneeded applications. This will ensure a better result with fewer false positives.

Sysinternals TCPView

TCPView is a program which shows the detailed listings of all TCP and UDP endpoints on your system, including the local and remote addresses and state of TCP connections. On Windows NT, 2000, and XP TCPView also reports the name of the process associated with the port.

Web page: www.microsoft.com/technet/sysinternals/utilities/TCPview.msp

You can download the program at the bottom of the page at the address above.

Version information: Freeware. Note that Microsoft now owns Sysinternals but still provides TCPview as a freeware program.

TCPView is more feature-rich and convenient to use than the NetStat program that comes with Windows. The .ZIP package also includes `Tcpvcon`, a command-line version of TCPView. This tool can be of use and help for detecting evidence of rootkit activity.

Copy the `TcpView.zip` file to any desired location on your hard drive. Extract the contents of this .ZIP archive to a dedicated folder, using the Extract All option on the context menu (Windows XP and later) or by using any third-party file-compression tool.

Double-click the `TcpView.exe` file to start the program. When you start the program, it enumerates all active TCP and UDP connections and resolves all IP addresses to their domain names. You can make TCPView to show IP addresses instead of domain names by de-selecting the Options ⇨ Resolve Addresses option. Whenever a new connection is established, TCPView displays it with a green background; when a connection is terminated, it's displayed with a red background. Connections that change their states are displayed in yellow. You can control the refresh rate of TCPView by selecting the desired time interval from the View ⇨ Update Speed menu.

UnHackMe



Your DART CD includes UnHackMe, one of the easiest tools to use for detecting and removing rootkits — completely automatically. It also has a real-time guard that can prevent the installation of rootkits.

Web site: www.greatis.com/unhackme

Version information: Trialware with a 30-day trial period. It's a commercial program, and it can be used freely for 30-days.

Double-click the name of the program on the CD to start the installation. During the installation, you'll be presented with an option to change the default install location. If you do not wish to change the location, continue the installation with all the default options.

To scan the system for rootkits, start UnHackMe and click the button Check. Next, check the button Check Me Now! If it finds any rootkits, it will show the details such as files, Registry entries, and services, etc. To automatically remove the rootkit, click the Stop button. UnHackMe automatically restarts the system and removes the detected rootkit on reboot.

Bonus Chapter 2

Ten (Plus Four) More Utilities

In This Chapter

- ▶ Acronis True Image
- ▶ Genie-Soft Backup Manager Home
- ▶ Karen's Replicator
- ▶ NTI Backup NOW!
- ▶ Sandboxie
- ▶ CCleaner
- ▶ Pocket KillBox
- ▶ Registrar Registry Manager
- ▶ Advanced Password Generator
- ▶ Any Password
- ▶ Password Safe
- ▶ Eraser
- ▶ KillDisk
- ▶ SpinRite

To continue the theme started in Bonus Chapter 1, just like soldiers on a real battlefield, you need to be ready for anything, even for defeat. Here we provide the tools and the know-how, not only to get you started but to carry you through the toughest and darkest hours.



You'll find some of these applications on your DART CD (for a complete list, see the Appendix); we've identified them with this icon. Others are easily available online; we include their Web sites in each entry.

Backup and Imaging Applications

Unlike some other forms of malware, rootkits can leave your files and applications damaged and unusable. Sometimes curing the disease can kill the patient, so it's prudent and wise to have clean, full backups of important files and programs, even your entire hard drive. These tools help you do that.

Acronis True Image



Your DART CD includes the trial version of Acronis True Image Home, an easy-to-use program for backing up selected files, folders, disks, or partitions. The backups can be stored on hard drives, CDs, Flash memory sticks, and other high-capacity media.

Web sites: Use one to read about it and the other to get it.

Overview: www.acronis.com/homecomputing/products/trueimage

Download: www.acronis.com/homecomputing/download/trueimage

Version information: Trialware has a 15-day trial period. You can download a user-guide PDF file, which has exhaustive information on all the features of Acronis True Image, from www.acronis.com/homecomputing/download/docs.



Be sure to check the system requirements in the Appendix or at the Acronis Web site before installing.

Double-click the program name on the CD to start the installation. In the install menu, select Acronis True Image as the program to install and follow the on-screen instructions given by the Installation wizard to complete the installation.

Start the program, either by double-clicking the desktop icon or by choosing the shortcut in the Start Menu. The main window of Acronis True Image offers options for handling tasks such as Backup, Recovery, Clone Disk, and so on. Each task is presented with an icon and a short description. To complete a task, click its icon and follow the instructions given by the wizard. (For example, to create a backup of a particular folder, click the Backup icon to open the Create Backup wizard, which guides you through the backup process.)

Genie-Soft Backup Manager Home

Genie Backup Manager Home is the one of the most user-friendly backup products on the market. Its wizard-based interface makes backing up an easy job.

Web site: www.genie-soft.com/products/gbm_shared/rnt.html

Version information: Direct purchase only.

Double-click the install file to start the installation. Follow the on-screen instructions to complete the installation. While installing, if needed, you can change the install location of the program.



Karen's Replicator

Your DART CD includes Karen's Replicator — a free, easy-to-use program for making scheduled automatic backups.

Web site: www.karenware.com/powertools/ptreplicator.asp

Version information: Freeware (for personal and noncommercial use). Obtain a paid commercial license if you want to use it at work. You can do so at www.karenware.com/sitelic.asp.

Double-click the program name on the CD to start the installation. You're presented with an option to make Replicator available to all the users in the system or only to yourself. Select the desired option and continue the installation, following the on-screen instructions.

NTI Backup NOW!

NTI Backup NOW! is a powerful backup program that can handle file, folder, or drive backups. The intuitive interface (called Easy Steps) makes all the backup jobs a breeze.

Web site: www.ntius.com

Version information: Trialware with a trial period of 30 days or seven uses. You can download it at the address below. A manual and tutorial are available for download at www.ntius.com/backup_now_downloads.asp.

Double-click the install file to start the installation. Read and accept the EULA agreement; specify an installation location if you want, and then follow the on-screen instructions displayed by the Installation wizard.

Sandboxie



Your DART CD includes Sandboxie — virtualization software that provides a sandbox-like environment for running programs. Sandboxie acts as a go-between, standing between your hard drive and the program you are using. Using Sandboxie to *sandbox* a Web browser in this way, it's possible to isolate your system from the threats of the Internet. Visit the Sandboxie Web site to see an illustrated explanation of what it does.



A *sandbox* is a special Java development environment with strict rules about what can be loaded to a computer with (in this case) a Web browser. It prevents Web pages designed with malicious intent from loading to the computer, using strict limits set on the computer resources requested or accessed. When surfing with sharks, it's better to stay in a sandbox. (So to speak.) When you sandbox your Web browser, it all exists within a single window. You are using a *virtual* Web browser rather than a real one. It still works the same but if you have any problems, such as malware trying to take hold of it, you simply turn it off and the malware is gone. The malware cannot write instructions to your hard drive because the sandbox prevents that. More information about virtual operating systems can be found in Chapter 6.

Web site: www.sandboxie.com

Version information: Freeware.

Double-click the name of the program on the CD to start the installation procedure. Follow the on-screen instructions to complete the installation. While installing Sandboxie, you may choose to change the installation location. Sandboxie places an icon in the System Tray, which you can use to open the program.

Right-click the Sandboxie icon to get the Sandboxie Control menu. You can use the options provided in this menu to launch any other program; to launch Internet Explorer (for example), simply click the Run Internet Explorer option. To launch any other program, click the Run Program option and then provide the path to the program you want to launch through Sandboxie. For a complete tutorial on using Sandboxie, visit this Web page:

www.sandboxie.com/usage.php

System and Registry Cleaners

How often do you take your car to the carwash? Did you know that it's just as important to keep the engine clean too? It's likewise with computers. Keeping the system clean and shiny makes it easier to fix when you have problems. Likewise with your Registry. (You can learn more about cleaning your system and Registry in Chapter 3.)

CCleaner

CCleaner searches for — and removes — junk files that may be present in the system. In addition to these file-cleaning operations, it can remove stray Registry entries.

Web site: www.ccleaner.com

Download: www.ccleaner.com/download

Version information: Freeware.

Double-click the install file to start the installation. By default CCleaner installs Yahoo Toolbar along with its own program files; if you don't want that toolbar installed, uncheck the option when you see it during the installation procedure.

Pocket KillBox

Pocket KillBox is a program for deleting stubborn or malware files that simply refuse to get deleted by normal methods.

Web site: www.bleepingcomputer.com/files/killbox.php

Version information: Freeware.

Copy the .ZIP archive to your hard drive. Extract the contents of the .ZIP archive to a dedicated folder, using the Extract All option present in the context menu (Windows XP and later). If you're using any third-party file-compression program, you can use it to extract the .ZIP archive.

To use the program, run `KillBox.exe`. Type the name of the file you want to delete, along with its path, in the *Full path of file to delete* text box. Then click the button that has a white cross on a red circle. If you'd prefer KillBox to delete the file at your next system boot-up, choose Delete On Reboot. (This option is especially useful for deleting rootkit files detected by Sysinternals Rootkit Revealer.)

Registrar Registry Manager

Registrar Registry Manager is an advanced Registry-management program. It has a wealth of features, including advanced search, Registry backup, and remote Registry maintenance. The Pro version is commercial software with additional capabilities such as defragmenting and cleaning the Registry.

Web site: www.resplendence.com/registrar

Version information: Free trial version (Lite Edition) and Pro Edition available. Trialware (command-line version) with a 21-day trial period.

Double-click the install file to start the installation. Accept the license agreement and proceed as per the on-screen instructions given by the Installation wizard to complete the installation.



Be careful while editing the Registry! Errors in the Registry may result in a corrupted — or irreparable — system. Be sure to set a fresh Restore point before using this utility. See Chapter 3 to learn how to set Restore points.

Password Protectors and Generators

Too many people set weak passwords or use none at all because they can't remember them well. The following sections describe some applications to assist you in generating and storing good passwords.

Advanced Password Generator



Your DART CD includes Advanced Password Generator — a program for generating secure passwords by using special algorithms called *random-number generators*. The free version can generate two passwords at a time, each four characters long. The commercial, registered version can generate up to 9999 passwords at a time, each up to 34 characters long.

Web site: www.segobit.com/apg.htm

Version information: Shareware.

Double-click the name of the program on the CD to start the installation. Accept the license agreement, choose the location where the program is to be installed, and follow the on-screen instructions to finish the installation.

Start the program, either by using its Desktop icon or from the Start Menu shortcut. First, select any one of the desired algorithms from the Random Number Generator option box — and then the mode of operation from the Mode option box. You can also choose the characters that are to be included in the password by selecting the appropriate option from the Keys option box. Next, you can specify whether the passwords should be case-sensitive by selecting the required option in the Case Sensitive or Not option box. Finally, you can select the password length and password quantities from the respective options shown on the main window; click Generate Password to generate passwords of the desired specifications.

Any Password



Your DART CD includes Any Password — a tool you can use to store all your passwords, user IDs, and related information in a tree-structured database — as encrypted, password-protected files. You need only remember the master password.

Web site: www.anypassword.com/index.html

Version information: Freeware.

Any Password is a personal favorite of ours. Unlike Password Safe (coming up in a minute), it allows you to copy and paste long passwords whenever you want to. If you're fussy about removing passwords from the Clipboard, simply close the program. Any Password will automatically clear the Clipboard for you.

Double-click the program name on the CD to start the installation. Follow the on-screen instructions, where you need to accept the license agreement and, if needed, choose the desired destination folder to complete the installation.

Launch the program either by using the Start Menu shortcut or its Desktop icon. To create a new database record that can store multiple passwords, choose the Edit ⇨ Add Folder option. This will create a new database record and will be displayed under the My Passwords tree in the left pane of Any Password program window. Now to store the passwords and related information in this new database, choose the Edit ⇨ New Password option, and fill

the various the fields given in the right pane of the program window, with the respective information. For example, to store the password of an e-mail account, type the e-mail address in the User Name/Account field, e-mail password in Password field, and finally the e-mail provider's Web site in the URL/File/Program field. You can type any additional information that may be needed in the Comments field.

Multiple passwords can be stored in a single record if you choose the Edit ⇄ Add Password option repeatedly. When you're done, save the changes made to the database by choosing File ⇄ Save option.

While saving, you'll be prompted to provide a master password for the file. To view the stored passwords, you can open this file by providing the master password.

Password Safe

Password Safe is a free program that securely stores all passwords in encrypted databases, each requiring a master password to open it. You can move these databases and use them in different systems, provided those systems are running the identical version of Password Safe.

Web site: <http://passwordsafe.sourceforge.net>

Download Mirrors (select one): <http://prdownloads.sourceforge.net/passwordsafe/pwsafe-3.04.exe?download>

Version information: Freeware.

Double-click the install file to start the installation. Accept the license agreement and proceed with the installation.

After completing installation, you can start the program by using its Desktop icon or Start Menu shortcut.

When you run the program for the first time, your first task is to create a database. To do this, click the Create new database button. Next, type the desired master password in the "Safe Combination" text box and retype it in the "Verify" text box and click OK. This opens the main window of Password Safe. To add passwords to the newly created record, click Edit ⇄ Add Entry. This will show the Edit/View Entry dialog box, where you provide information such as username, password, group to which the account belongs, and so on. When you've filled the needed fields, click OK to save the entry in the database. You can add multiple password entries by selecting the Edit ⇄ Add Entry option again. After adding the required number of password entries to the database, save the database by choosing File ⇄ Save. Thereafter, to open the databases you've saved, you have to provide the master password for each one.

Password Safe can also create backups of database records. To do so, you choose **Manage** ⇨ **Make Backup** and type in the path to the database file that you want backed up.



If you lose the master password of a database, then there is no way you can recover the password accounts stored in that database. So don't forget the master password! (But you knew that.)

Hard-Drive Erase and Repair Utilities

You may win the battles against the rootkit but the patient might not be saved. Your computer may have been rendered unusable by damaged and corrupted system files, or it may have been invaded by a rootkit that hides itself in bad sectors on the hard drive. What you can do about it is presented here. How these programs are used together is discussed in Chapter 11.

Eraser

Eraser is an advanced, freeware security tool that allows you to completely remove data from hard drives. Eraser destroys data by overwriting it several times with patterns you select.

Web site: www.heidi.ie/eraser

Version information: Freeware.

Download the Eraser installation file (`Eraser57Setup.zip`) from www.heidi.ie/eraser/download.php. Save it to a convenient location on your hard drive. Extract the contents of this .ZIP archive to a dedicated folder, using the **Extract All** option in the context menu (Windows XP and later) or by using a third-party file-compression tool. Now, start the installation of Eraser by double-clicking the `EraserSetup.exe` file. Follow the on-screen instructions to complete the installation.

Eraser integrates itself with Windows Explorer; you can use Eraser to permanently erase a file or folder simply by right-clicking it and choosing **Erase** from the context menu.

When you click the **Erase** option, you're prompted with a dialog box that asks you to confirm the file or folder you're erasing, and you can change the erasing methods (by clicking the **Options** button) if you want. After making the desired erase-method selection, click **Yes** to permanently delete the file or folder.

Eraser can also be used to wipe complete drives from a bootable floppy or CD. To create a bootable floppy, use the Create Boot Nuke Disk option from the Start Menu program group in Eraser. Clicking this shortcut opens a program called WinImage Self Extractor. Here, select Writing on Floppy ⇔ Formatting. Insert a floppy disk into the appropriate drive and click OK. This creates a bootable floppy with Eraser, which you can use to wipe hard drives. If you don't have a floppy-disk drive, you can download the ISO bootable CD image file (named `dban-1.0.6_i386.iso`) from this link:

```
http://prdownloads.sourceforge.net/dban/dban-1.0.6_i386.iso?download
```

Then you can burn this ISO image to a CD to create a bootable Eraser CD. (If you need a refresher on how to burn ISO images to CD, you can find it in the Appendix.)

Active@ Kill Disk

Active@ Kill Disk is a DOS-based, small (but powerful) hard-drive-eraser program that allows you to wipe all the data on hard drives and floppy disks completely, without any possibility of data recovery.

Web site: www.killdisk.com

Version information: Freeware version

This program has to be run from a bootable medium such as a CD, floppy (yes, it will fit), or USB drive. You can use the ISO bootable CD image to create a bootable CD that already contains the KillDisk program. Here's the drill:

- 1. Download the zipped ISO image file, `boot-cd-iso.zip`, from the following link:**

```
http://download2.lsoft.net/boot-cd-iso.zip
```

- 2. Save the downloaded ISO file to a convenient location on your hard drive.**

- 3. Extract the contents of this .ZIP archive to a dedicated folder.**

You can do so by, using the Extract All option found on the context menu (Windows XP and later) or by using a third-party file-compression program. When you're done, you have an ISO image file.

- 4. Burn the ISO image file to CD and boot your PC from that disc.**



Here's where to find how to do that:

- For pointers on burning this ISO image file to a CD, refer to the Appendix.
- To Make your PC boot from the Active@ Kill Disk bootable CD you create, you need to make some changes in the BIOS — *very carefully*. Chapter 11 walks you through that process.

5. Insert the CD to the CD ROM drive and restart the PC.

Since you've already specified the CD-ROM drive as the first boot device, the Active@ Kill Disk CD should start automatically.

6. To run the program, type KILLDISK.EXE and press Enter.

You should now see the KillDisk user interface, showing all hard-drive partitions in the system.

7. Choose the partition(s) to wipe. To view the partitions press Ctrl+S.

Use the arrow keys to navigate through the hard-drive partition list. To wipe a partition, select that partition and press Enter.

KillDisk shows the Erase Method configuration window where you can select any one of the methods to be used to erase the disk.

8. Select an erase method, and press Enter to start the operation.

The time taken to complete the operation depends on the selected erase method.

To learn about all the options present in KillDisk, you can check the User Guide at their Web site, or you can refer to the PDF user guide available via this link:

www.killdisk.com/downloads/killdisk.pdf

SpinRite

SpinRite is a uniquely powerful program for preventing hard-drive crashes and recovering after a malware attack.

Web site: www.grc.com/sr/spinrite.htm

Version information: Direct purchase only. SpinRite is commercial software; it doesn't have a trial version. You can visit the SpinRite Web page to buy it.

Running SpinRite regularly prevents disk crashes. Any weak and failing areas within the regions on the hard drive are identified and removed from use. If SpinRite is used for crash recovery, it uses special techniques and algorithms to restore as much data as possible from bad sectors of a crashed disk.

SpinRite should be run either from a bootable CD or a floppy. Start the SpinRite program and click the Create ISO or IMG file button. Another window is displayed that has instructions for creating an ISO file. Here, click the Save a Boot Image File button. SpinRite will now create a bootable ISO image file called `SpinRite.iso`. This ISO image file should be recorded to a blank CD (the Appendix sums up the steps for doing this).

In order to use SpinRite, you have to boot the PC from this prepared CD-ROM. To find out how to make a PC boot from CD, refer to Chapter 11 — and follow its instructions for editing the BIOS very carefully. (Okay, it sounds like nagging — but you'll thank us later.)

Insert the SpinRite CD into the CD-ROM drive and restart the PC. The PC boots from the SpinRite CD and displays a command prompt. To start SpinRite, type **SPINRITE** at the command prompt and press Enter. SpinRite displays a menu, from which you choose the Begin SpinRite option. SpinRite displays all available hard drives and partitions. Here, after you select your desired partitions, SpinRite analyzes them and displays their status. To recover data from a crashed disk, you can select the DynaStat Data Recovery option from the menu that's always displayed on-screen; then you follow the on-screen instructions to start recovery. To explore all the options SpinRite offers, you can refer to a handy PDF user guide available at this link:

www.grc.com/srdocs.htm