

Inside Windows® Debugging

Practical Debugging
and Tracing Strategies



Tarik Soulami

Inside Windows® Debugging: Practical Debugging and Tracing Strategies

Use Windows debuggers throughout the development cycle—and build better software

Rethink your use of Windows debugging and tracing tools—and learn how to make them a key part of test-driven software development. Led by a member of the Windows Fundamentals Team at Microsoft, you'll apply expert debugging and tracing techniques—and sharpen your C++ and C# code analysis skills—through practical examples and common scenarios. Learn why experienced developers use debuggers in every step of the development process, and not just when bugs appear.

Discover how to:

- Go behind the scenes to examine how powerful Windows debuggers work
- Catch bugs early in the development cycle with static and runtime analysis tools
- Gain practical strategies to tackle the most common code defects
- Apply expert tricks to handle user-mode and kernel-mode debugging tasks
- Implement postmortem techniques such as JIT and dump debugging
- Debug the concurrency and security aspects of your software
- Use debuggers to analyze interactions between your code and the operating system
- Analyze software behavior with the Event Tracing for Windows (ETW) framework



Get code samples on the web

Ready to download at
<http://go.microsoft.com/fwlink/?Linkid=245713>

For **system requirements**, see the Introduction.

microsoft.com/mspress

ISBN: 978-0-7356-6278-0



9 780735 662780

U.S.A. \$39.99

Canada \$41.99

[Recommended]

Programming/Windows Debugging



About the Author

Tarik Soulami, a principal development lead on the Windows Fundamentals Team, has more than 10 years of experience designing and developing system-level software at Microsoft. Before joining the Windows team, Tarik spent several years on the Common Language Runtime (CLR) Team, where he helped shape early releases of the Microsoft® .NET framework.

DEVELOPER ROADMAP

Start Here!

- Beginner-level instruction
- Easy to follow explanations and examples
- Exercises to build your first projects



Step by Step

- For experienced developers learning a new topic
- Focus on fundamental techniques and tools
- Hands-on tutorial with practice files plus eBook



Developer Reference

- Professional developers; intermediate to advanced
- Expertly covers essential topics and techniques
- Features extensive, adaptable code examples



Focused Topics

- For programmers who develop complex or advanced solutions
- Specialized topics; narrow focus; deep coverage
- Features extensive, adaptable code examples



Microsoft®

Inside Windows[®] Debugging

Tarik Soulami

Copyright © 2012 by Tarik Soulami

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

ISBN: 978-0-7356-6278-0

2 3 4 5 6 7 8 9 10 LSI 8 7 6 5 4 3

Printed and bound in the United States of America.

Microsoft Press books are available through booksellers and distributors worldwide. If you need support related to this book, email Microsoft Press Book Support at mspinput@microsoft.com. Please tell us what you think of this book at <http://www.microsoft.com/learning/booksurvey>.

Microsoft and the trademarks listed at <http://www.microsoft.com/about/legal/en/us/IntellectualProperty/Trademarks/EN-US.aspx> are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

The example companies, organizations, products, domain names, email addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred.

This book expresses the author's views and opinions. The information contained in this book is provided without any express, statutory, or implied warranties. Neither the authors, Microsoft Corporation, nor its resellers, or distributors will be held liable for any damages caused or alleged to be caused either directly or indirectly by this book.

Acquisitions Editor: Russell Jones

Developmental Editor: Russell Jones

Production Editor: Melanie Yarbrough

Editorial Production: Waypoint Press

Technical Reviewer: John Mueller

Copyeditor: Roger LeBlanc

Indexer: Christina Yeager

Cover Design: Twist Creative • Seattle

Cover Composition: Karen Montgomery

Illustrator: Steve Sagman

Contents at a Glance

	<i>Foreword</i>	xv
	<i>Introduction</i>	xvii
<hr/>		
PART I	A BIT OF BACKGROUND	
<hr/>		
CHAPTER 1	Software Development in Windows	3
<hr/>		
PART II	DEBUGGING FOR FUN AND PROFIT	
<hr/>		
CHAPTER 2	Getting Started	33
CHAPTER 3	How Windows Debuggers Work	85
CHAPTER 4	Postmortem Debugging	125
CHAPTER 5	Beyond the Basics	159
CHAPTER 6	Code Analysis Tools	195
CHAPTER 7	Expert Debugging Tricks	219
CHAPTER 8	Common Debugging Scenarios, Part 1	267
CHAPTER 9	Common Debugging Scenarios, Part 2	323
CHAPTER 10	Debugging System Internals	365
<hr/>		
PART III	OBSERVING AND ANALYZING SOFTWARE BEHAVIOR	
<hr/>		
CHAPTER 11	Introducing Xperf	391
CHAPTER 12	Inside ETW	415
CHAPTER 13	Common Tracing Scenarios	457
APPENDIX A	WinDbg User-Mode Debugging Quick Start	505
APPENDIX B	WinDbg Kernel-Mode Debugging Quick Start	519
	<i>Index</i>	527

Contents

<i>Foreword</i>	xv
<i>Introduction</i>	xvii
Who Should Read This Book	xviii
Assumptions	xviii
Organization of This Book	xviii
Conventions in This Book	xix
System Requirements	xx
Code Samples	xxi
Installing the Code Samples	xxi
Running the Code Samples	xxii
Acknowledgments	xxvi
Errata & Book Support	xxvii
We Want to Hear from You	xxvii
Stay in Touch	xxviii

PART I A BIT OF BACKGROUND

Chapter 1 Software Development in Windows	3
Windows Evolution	3
Windows Release History	3
Supported CPU Architectures	4
Windows Build Flavors	5
Windows Servicing Terminology	6
Windows Architecture	7
Kernel Mode vs. User Mode	8
User-Mode System Processes	9
User-Mode Application Processes	10
Low-Level Windows Communication Mechanisms	13
Windows Developer Interface	16

Developer Documentation Resources	16
WDM, KMDf, and UMDF	17
The NTDLL and USER32 Layers	18
The Win32 API Layer	18
The COM Layer	19
The CLR (.NET) Layer	25
Microsoft Developer Tools	28
The Windows DDK (WDK)	29
The Windows SDK	29
Summary	30

PART II DEBUGGING FOR FUN AND PROFIT

Chapter 2 Getting Started **33**

Introducing the Debugging Tools	34
Acquiring the Windows Debuggers Package	34
Acquiring the Visual Studio Debugger	38
Comparing the WinDbg and Visual Studio Debuggers	38
User-Mode Debugging	39
Debugging Your First Program with WinDbg	39
Listing the Values of Local Variables and Function Parameters	47
Source-Level Debugging in WinDbg	52
Symbol Files, Servers, and Local Caches	53
Caching Symbols Offline for WinDbg	55
Troubleshooting Symbol Resolution Issues in WinDbg	56
Name Decoration Considerations	57
Getting Help for WinDbg Commands	58
Kernel-Mode Debugging	60
Your First (Live) Kernel Debugging Session	61
Setting Up a Kernel-Mode Debugging Environment Using Physical Machines	67
Setting Up a Kernel-Mode Debugging Environment Using Virtual Machines	73

Diagnosing Host/Target Communication Issues	76
Understanding the KD Break-in Sequence	77
Controlling the Target in the Kernel Debugger	78
Setting Code Breakpoints in the Kernel Debugger	81
Getting Help for WinDbg Kernel Debugging Commands	83
Summary.	83

Chapter 3 How Windows Debuggers Work 85

User-Mode Debugging	85
Architecture Overview.	86
Win32 Debugging APIs.	87
Debug Events and Exceptions	88
The Break-in Sequence	91
Setting Code Breakpoints.	93
Observing Code Breakpoint Insertion in WinDbg	93
Kernel-Mode Debugging	98
Architecture Overview.	98
Setting Code Breakpoints.	100
Single-Stepping the Target.	100
Switching the Current Process Context	101
Managed-Code Debugging	103
Architecture Overview.	103
The SOS Windows Debuggers Extension.	106
Script Debugging	112
Architecture Overview.	112
Debugging Scripts in Visual Studio.	114
Remote Debugging	116
Architecture Overview.	116
Remote Debugging in WinDbg.	117
Remote Debugging in Visual Studio.	121
Summary.	123

Chapter 4	Postmortem Debugging	125
	Just-in-Time Debugging	125
	Your First JIT Debugging Experiment	126
	How Just-in-Time Debugging Works	128
	Using Visual Studio as Your JIT Debugger	132
	Run-Time Assertions and JIT Debugging	138
	JIT Debugging in Session 0	139
	Dump Debugging	139
	Automatic User-Mode, Crash-Dump Generation	139
	Analyzing Crash Dumps Using the WinDbg Debugger	143
	Analyzing Crash Dumps in Visual Studio	150
	Manual Dump-File Generation	151
	“Time Travel” Debugging	153
	Kernel-Mode Postmortem Debugging	153
	Summary	157
Chapter 5	Beyond the Basics	159
	Noninvasive Debugging	159
	Data Breakpoints	162
	Deep Inside User-Mode and Kernel-Mode Data Breakpoints	163
	Clearing Kernel-Mode Data Breakpoints	165
	Execution Data Breakpoints vs. Code Breakpoints	166
	User-Mode Debugger Data Breakpoints in Action: C++ Global Objects and the C Runtime Library	168
	Kernel-Mode Debugger Data Breakpoints in Action: Waiting for a Process to Exit	170
	Advanced Example: Who Is Changing a Registry Value?	172
	Scripting the Debugger	176
	Replaying Commands Using Debugger Scripts	176
	Debugger Pseudo-Registers	178
	Resolving C++ Template Names in Debugger Scripts	180
	Scripts in Action: Listing Windows Service Processes in the Kernel Debugger	181

WOW64 Debugging	183
The WOW64 Environment	183
Debugging of WOW64 Processes	184
Windows Debugging Hooks (GFLAGS)	187
Systemwide vs. Process-Specific NT Global Flags	187
The GFLAGS Tool	188
The <i>!gflag</i> Debugger Extension Command	191
Impact of User-Mode Debuggers on the Value of the NT Global Flag	193
The Image File Execution Options Hooks	193
Summary	194

Chapter 6 Code Analysis Tools 195

Static Code Analysis	195
Catching Your First Crashing Bug Using VC++ Static Code Analysis	196
SAL Annotations	199
Other Static Analysis Tools	202
Runtime Code Analysis	206
Catching Your First Bug Using the Application Verifier Tool	206
A Behind-the-Scenes Look: Verifier Support in the Operating System	209
The <i>!avrf</i> Debugger Extension Command	214
The Application Verifier as a Quality Assurance Tool	217
Summary	217

Chapter 7 Expert Debugging Tricks 219

Essential Tricks	220
Waiting for a Debugger to Attach to the Target	220
Breaking on DLL Load Events	222
Debugging Process Startup	227
Debugging Child Processes	234

More Useful Tricks	245
Debugging Error-Code Failures.....	245
Breaking on First-Chance Exception Notifications.....	252
Freezing Threads	253
Kernel-Mode Debugging Tricks.....	255
Breaking on User-Mode Process Creation.....	255
Debugging the Startup of User-Mode Processes.....	259
Breaking on DLL Load Events.....	260
Breaking on Unhandled SEH Exceptions.....	262
Freezing Threads	262
Summary.....	265

Chapter 8 Common Debugging Scenarios, Part 1 267

Debugging Access Violations	267
Understanding Memory Access Violations	268
The <i>!analyze</i> Debugger Extension Command	269
Debugging Heap Corruptions	271
Debugging Native Heap Corruptions.....	271
Debugging Managed (GC) Heap Corruptions	281
Debugging Stack Corruptions	291
Stack-Based Buffer Overruns	291
Using Data Breakpoints in Stack Corruption Investigations.....	294
Reconstructing Call Frames from Corrupted Stacks	295
Debugging Stack Overflows	297
Understanding Stack Overflows	297
The <i>kf</i> Debugger Command	299
Debugging Handle Leaks	300
A Handle Leak Example	300
The <i>!htrace</i> Debugger Extension Command.....	302
Debugging User-Mode Memory Leaks	307
Detecting Resource Leaks Using the Application Verifier Tool ..	307
Investigating Memory Leaks Using the UMDH Tool	310
Extending the Strategy:	
A Custom Reference Stack-Trace Database.....	314

Debugging Kernel-Mode Memory Leaks	316
Kernel Memory Basics	316
Investigating Kernel-Mode Leaks Using Pool Tagging	318
Summary.	322

Chapter 9 Common Debugging Scenarios, Part 2 323

Debugging Race Conditions	323
Shared-State Consistency Bugs	324
Shared-State Lifetime Management Bugs	330
DLL Module Lifetime-Management Bugs	340
Debugging Deadlocks	343
Lock-Ordering Deadlocks	344
Logical Deadlocks	348
Debugging Access-Check Problems	352
The Basic NT Security Model	353
Windows Vista Improvements	358
Wrapping Up.	362
Summary.	363

Chapter 10 Debugging System Internals 365

The Windows Console Subsystem	365
The Magic Behind <i>printf</i>	366
Handling of Windows UI Events	373
Handling of the Ctrl+C Signal	374
Anatomy of System Calls	380
The User-Mode Side of System Calls	381
The Transition into Kernel Mode	383
The Kernel-Mode Side of System Calls.	385
Summary	387

Chapter 11 Introducing Xperf	391
Acquiring Xperf	391
Your First Xperf Investigation	396
Devising an Investigation Strategy	397
Collecting an ETW Trace for the Scenario	397
Analyzing the Collected ETW Trace	399
Xperf's Strengths and Limitations	411
Summary	412
Chapter 12 Inside ETW	415
ETW Architecture	416
ETW Design Principles	416
ETW Components	417
The Special NT Kernel Logger Session	418
Configuring ETW Sessions Using Xperf	419
Existing ETW Instrumentation in Windows	422
Instrumentation in the Windows Kernel	422
Instrumentation in Other Windows Components	426
Understanding ETW Stack-Walk Events	431
Enabling and Viewing Stack Traces for Kernel Provider Events	432
Enabling and Viewing Stack Traces for User Provider Events	434
Diagnosing ETW Stack-Trace Issues	436
Adding ETW Logging to Your Code	441
Anatomy of ETW Events	441
Logging Events Using the ETW Win32 APIs	445
Boot Tracing in ETW	449
Logging Kernel Provider Events During Boot	450
Logging User Provider Events During Boot	452
Summary	455

Chapter 13 Common Tracing Scenarios 457

Analyzing Blocked Time	458
The <i>CSwitch</i> and <i>ReadyThread</i> ETW Events	459
Wait Analysis Using Visual Studio 2010	461
Wait Analysis Using Xperf	467
Analyzing Memory Usage	473
Analyzing High-Level Memory Usage in a Target Process	474
Analyzing NT Heap Memory Usage	475
Analyzing GC Heap (.NET) Memory Usage	481
Tracing as a Debugging Aid	490
Tracing Error Code Failures	490
Tracing System Internals	494
Summary	502

Chapter 14 WinDbg User-Mode Debugging Quick Start 505

Starting a User-Mode Debugging Session	505
Fixing the Symbols Path	505
Fixing the Sources Path	506
Displaying the Command Line of the Target Process	507
Control Flow Commands	507
Listing Loaded Modules and Their Version	508
Resolving Function Addresses	509
Setting Code (Software) Breakpoints	509
Setting Data (Hardware) Breakpoints	510
Switching Between Threads	511
Displaying Call Stacks	511
Displaying Function Parameters	512
Displaying Local Variables	513
Displaying Data Members of Native Types	513

Navigating Between Call Frames	514
Listing Function Disassembly	515
Displaying and Modifying Memory and Register Values	516
Ending a User-Mode Debugging Session.	518
Chapter 15 WinDbg Kernel-Mode Debugging Quick Start	519
Starting a Kernel-Mode Debugging Session	519
Switching Between CPU Contexts	519
Displaying Process Information	520
Displaying Thread Information.	521
Switching Process and Thread Contexts	522
Listing Loaded Modules and Their Version	523
Setting Code (Software) Breakpoints Inside Kernel-Mode Code.	524
Setting Code (Software) Breakpoints Inside User-Mode Code.	525
Setting Data (Hardware) Breakpoints	525
Ending a Kernel-Mode Debugging Session	526
<i>Index</i>	527
<i>About the Author</i>	561

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

microsoft.com/learning/booksurvey

Foreword

Like many others, I am a firm believer in using tools to expand our understanding of how systems really work. In fact, I began my career as a performance tools developer. My boss at that time had many simple sayings; among them was one of my favorites, “Our team is only as good as our people and our people are only as good as our tools.” As a manager, I’ve made it a priority to ensure we dedicate many of our top engineers to tools development and I have always encouraged using tools to teach and grow engineers.

Teaching books such as Tarik’s are meant to help improve the productivity, understanding, and confidence of others. Being an individual who has had both a lifelong passion for tools and learning, as well as someone who has spent more than a decade and a half working in Windows, it is an honor and a pleasure for me to write the foreword to Tarik’s insightful book.

For decades, the Windows team has worked tirelessly to improve the core capabilities of the platform and to make it more suitable for increasingly diverse hardware configurations and software stacks. This hard work has paid off; Windows today is the preeminent platform for developers, consumers, and businesses across the globe. With more than 1 billion PCs and customers, Windows is both the market-leading server and client computing platform. Variants of Windows run on small form-factor mobile devices, within embedded intelligent systems, on uniform and non-uniform (NUMA) memory architectures, on the XBOX gaming console, and on single CPU systems as well as those with hundreds of processors. Windows runs in clustered configurations with failover capabilities and sits atop hypervisors and Virtual Machines. And of course, Windows is in the cloud.

Enormous platform success and diversity can be accompanied by an equally enormous and diverse set of technical challenges. Thousands of engineers at Microsoft, and tens to hundreds of thousands outside of Microsoft, are involved in building, debugging, and troubleshooting a multitude of diverse configurations and solutions. To understand issues in a vastly diverse problem space, foundational tools, techniques, and concepts are essential.

To that end, Tarik has done an excellent job in detailing how a set of key foundational Windows tools work. In detailing these tools, Tarik succeeds in expanding the reader's awareness of key operating system concepts, architectures, strengths, and limitations. With the concepts understood and the tools mastered, readers can expect to be able to tackle all types of performance and debugging challenges. I believe Tarik's book will become a staple for a broad set of individuals, from novices to experts.

*Michael Fortin, Ph.D.
Distinguished Engineer, Windows Fundamentals
Microsoft Corporation*

Introduction

One exciting aspect of software programming is that there are usually many ways to accomplish the same goal. Unfortunately, this also presents software engineers with unique challenges when trying to make the best design or implementation choice for each situation. Experience plays a major role, and the learning process is often progressive as one learns to avoid past mistakes. Sadly, though, experience is often a variable concept. I have met several software engineers who, after spending a very long time working on an area, still lacked a basic understanding of how it *really* worked beyond the repetitive day-to-day tasks they grew accustomed to. I have also met others who have perfected their craft in a field after only a few years of working experience.

This book introduces a few techniques for methodically approaching software development problems primarily using the two “Swiss Army knives” of expert Microsoft Windows developers—namely, the Windows debuggers (WinDbg) and the Windows Performance Toolkit (Xperf). By focusing on *why* features and components work the way they do in the system rather than simply on *how* they work or *what* they do, this book tries to accelerate the process of learning by experience and to minimize the number of mistakes made when approaching new problems. An important part of the process is learning to compare and contrast with known solutions to existing ones.

Software engineering is still inherently a practical science. (Some might even argue it’s an art rather than a science.) While there is certainly no substitute for real experience, the topic can definitely be approached with the same methodical persistence that works so well for other scientific disciplines. In fact, this approach works even better in software engineering because all behaviors can be explained rationally. After all, it is all just code—whether it’s your own code or code written by others that you end up consuming in your software—and code can always be traced and understood.

Although this book deals with several architectural pillars of the Windows operating system as part of its debugging and tracing experiments, my main goal in writing it is less about covering those details and more about encouraging and developing this critical mindset. I hope to demonstrate how this approach can be used for solving a few interesting problems as part of this book, and that you continue to systematically apply debugging and tracing as you expand your learning beyond the topics directly covered here.

This book is not really about teaching native or managed code programming, either, although you’ll find several good coding examples in the companion source code. Because it takes an inside-out look at how to explore the system using debugging

and tracing tools, this book will probably appeal more to software engineers with the desire to understand system internals rather than those with a need to quickly learn how to make use of a specific technology. However, I believe the approach and mindset it aspires to inculcate are applicable regardless of the technology or level of expertise. In fact, contrary to what many think, the higher the level of technology involved, the harder it becomes to grasp what goes on behind the scenes and the more expertise is needed in order to investigate failures when things inevitably go awry and require debugging skills to save the day. In pure C, for example, a call to *malloc* is just that: a function call. In C++, a call to the *new* keyword is emitted by the compiler as a call to the *new* operator function to allocate memory for the object, followed by code to construct the object (again, emitted by the compiler to possibly initialize a virtual pointer and invoke the constructors of the base classes, construct member data objects, and finally invoke the user-provided constructor code for the target leaf class). In C# (.NET), things get even more involved, because a one-line call to the *new* keyword might involve compiling new code at runtime, performing security checks by the .NET execution engine, loading the modules where the target type is defined, tracking the object reference for later garbage collection, and so on.

Who Should Read This Book

This book is aimed at software engineers who desire to take their game to the next level, so to speak, and perfect their mastery of Windows as a development platform through the use of debugging and tracing tools.

Assumptions

Readers should have basic familiarity with the C/C++ and C# programming languages. A basic knowledge of the Win32 and .NET platforms is helpful, but not strictly required, because this book makes every effort to introduce basic concepts before expanding into more advanced topics.

Organization of This Book

This book is divided into three parts:

- Part 1, “A Bit of Background,” provides a brief overview of Windows development frameworks and the layers in the operating system that support them. This

basic knowledge is important to have when trying to make sense of the data surfaced by debugging and tracing tools.

- Part 2, “Debugging for Fun and Profit,” covers the architectural foundations of debuggers in the Windows operating system. It also presents a number of extensible strategies that will help you make the most of the Windows debuggers. In addition, this part also shows how to use the WinDbg debugger to better understand system internals by analyzing the important interactions between your code and the operating system.
- Part 3, “Observing and Analyzing Software Behavior,” continues this theme. It presents the Event Tracing for Windows (ETW) technology and illustrates how to leverage it in debugging and profiling investigations.
- Finally, you’ll find two short appendices at the end of the book that recap the most common debugging tasks and how to accomplish them using WinDbg.

The table of contents will help you locate chapters and sections quickly. In addition, each chapter starts with a list of the main points covered in the chapter, and concludes with a summary section. You can also download the source code for all the experiments and examples shown throughout the book.

Conventions in This Book

This book presents information using conventions (listed in the following table) designed to make the information readable and easy to follow:

Convention	Meaning
Sidebars	Boxed sidebars feature additional bits of information that might be helpful for a given subject.
Notes	Notes provide useful observations related to the content discussed in the main text.
Inline Code	Inline code—that is, code that appears within a paragraph—is shown in <i>italic</i> font.
Code Blocks	Code blocks are shown in a different font to help you distinguish code from text easily. Important statements appear in bold font to help you focus on those aspects of the code.
Debugger Listings	Debugger listings are shown in a different font, and important commands are bolded to highlight them. The listings are also often prefixed with the output from the standard <i>vertarget</i> debugger command, which displays the OS version and CPU architecture where the experiment was conducted.
Function Names	Function names are sometimes referenced using their WinDbg symbolic names. For example, <i>kernel32!CreateFileW</i> refers to the CreateFileW Win32 API (“W” for the Unicode flavor) exported by the <i>kernel32.dll</i> module.

System Requirements

You will need the following hardware and software to follow the experiments and code samples in this book:

- **Operating System** Windows Vista or later. Windows 7 (or Windows Server 2008 R2) is highly recommended.
- **Hardware** Any computer that supports the Windows 7 operating system (OS) requirements. Except for the live kernel debugging experiments, a second computer to serve as a host kernel-mode debugger machine is typically required for kernel debugging.



Note The target and host don't really need to be separate physical machines. A common kernel debugging configuration—detailed in Chapter 2, “Getting Started”—is to run the target machine where you conduct the experiments as a Windows 7 virtual OS on a Windows Server 2008 R2 physical host computer, and run the kernel debugger in the host OS.

- **Hard disk** 1 GB of free hard-disk space to download and save the Windows Software Development Kit (SDK) and Driver Development Kit (DDK) ISO images. 40 MB of free hard-disk space to download and compile the companion source code. An additional 3 GB is required for installing Microsoft Visual Studio 2010.
- **Software** The following tools are used in the debugging and tracing examples shown throughout this book:
 - Version 7.1 of the Windows 7 SDK, which can be downloaded from the Microsoft Download Center at <http://www.microsoft.com/download/en/details.aspx?id=8442>. Both the Windows Debuggers (WinDbg) and Windows Performance Toolkit (Xperf) are part of this SDK.
 - The Application Verifier tool, which can also be downloaded from the Microsoft Download Center at <http://www.microsoft.com/download/en/details.aspx?id=20028>.
 - The System Internals suite of developer tools, which can be downloaded from <http://technet.microsoft.com/en-us/sysinternals/bb842062>.
 - Visual Studio 2010, any edition (excluding the free, stripped-down Express edition). The Visual Studio Ultimate edition is preferred because some advanced features, such as static code analysis and performance profil-

ing, are not supported in other editions. The 90-day, free trial version of Visual Studio offered by Microsoft, which can be downloaded at <http://www.microsoft.com/download/en/details.aspx?id=12187>, should suffice.

- The Windows 7 Driver Development Kit (DDK) is used to compile the companion source code and can also be downloaded from the Microsoft Download Center at <http://www.microsoft.com/download/en/details.aspx?id=11800>.

Code Samples

Most of the chapters in this book include experiments and examples that let you interactively try out new material introduced in the main text. The programs used in these experiments can be downloaded from the following page:

<http://www.microsoftpressstore.com/title/9780735662780>

Follow the instructions to download the `Inside_Windows_Debugging_Samples.zip` file.

Installing the Code Samples

Follow these steps to install the code samples on your computer so that you can use them with the experiments and examples provided in this book:

1. Unzip the `Inside_Windows_Debugging_Samples.zip` file that you downloaded from the book's website into a folder named `\Book\Code`.



Warning Do not use directory names with spaces in the path hierarchy that you choose to download the samples to. The DDK build environment, which will be described shortly in the “Compiling the Code Samples” section, will fail to compile the source code if you do. It is recommended that you use `\Book\Code` as the root of the source code because this is the assumed location used when referencing the programs in the main text. The samples in the companion source code are organized by chapter, so there should be a folder for every chapter under this root path location.

2. If prompted, review the displayed end user license agreement. If you accept the terms, select the accept option, and then click Next.

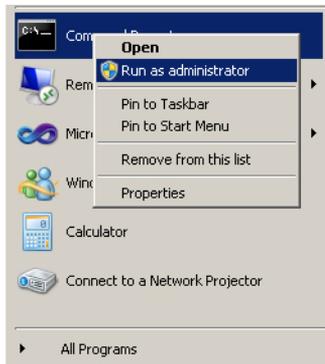


Note If the license agreement doesn't appear, you can access it from the same webpage from which you downloaded the `Inside_Windows_Debugging_Samples.zip` file.

Running the Code Samples

The code samples are organized by chapter and are referenced in the book main text by their respective directory path locations to help you find them easily.

Local Administrator rights are required by some sample programs. In Windows Vista and later, a command prompt must be launched as elevated in order to have full administrative privileges even when the user account is a member of the local built-in Administrators security group. To do so in Windows 7, for example, you need to right-click the Command Prompt menu item in the Windows Start menu, and select Run As Administrator, as shown in the following screen shot:



Compiling the Code Samples

The supporting programs used in the experiments presented in the book main text fall into three categories:

- **C++ samples** The binaries for these programs are deliberately omitted from the downloadable ZIP file. When you compile these code samples locally, WinDbg locates their symbols and source code files automatically. So the experiments presented in the early chapters of this book will “just work” in this configuration, without needing to specify the source and symbol locations explicitly in WinDbg. The steps to compile all the native C++ code samples at once are included later in this section.

- **C# (.NET) samples** For convenience, the compiled .NET programs are included in the downloadable ZIP file. You can use them as-is and you won't lose much, given WinDbg doesn't support source-level .NET debugging, but you can also follow the instructions included in the following section to recompile them if you prefer.
- **JavaScript and Visual Basic samples** These scripts are interpreted by the corresponding scripting engines and do not require compilation.

Compiling the .NET Code Samples

Compiling the .NET code samples from the companion source code requires version 4.0 of the Microsoft .NET Framework or later. Though this version isn't installed by default on Windows 7, .NET Framework 4.0 gets installed by many other dependent programs, such as Visual Studio 2010. You can also download and install a standalone version from the Microsoft download center at <http://www.microsoft.com/download/en/details.aspx?id=17851>.

Each C# code sample from the companion source code will have a helper compilation script under the same directory. This script uses the .NET 4.0 C# compiler directly and is easy to invoke, as illustrated in the following command:

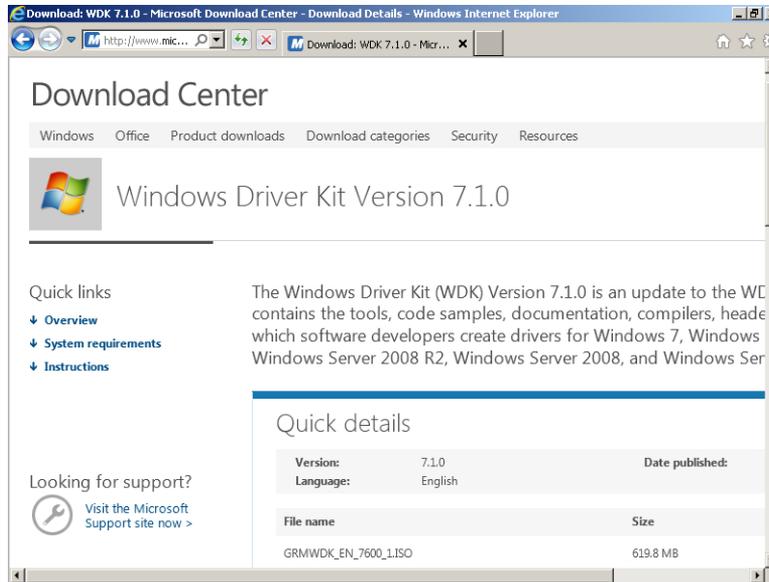
```
C:\book\code\chapter_04\LoadException>compile.bat
```

If this script fails to find the C# compiler, you should verify that you've installed .NET 4.0 in the default directory location where this script expects it to exist. If .NET 4.0 exists but was installed to a different location on your system, you need to modify the script and provide that path instead.

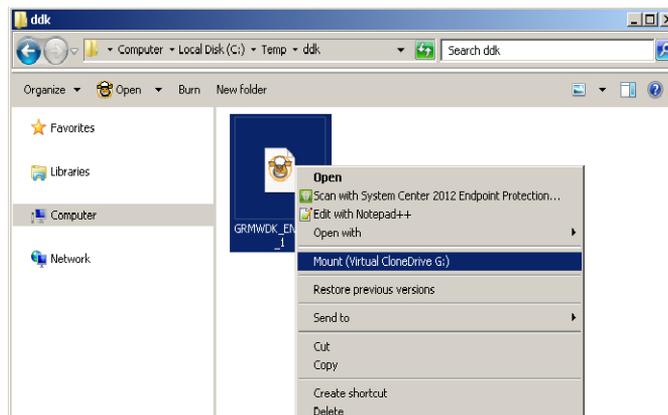
Compiling the C/C++ Code Samples

You can compile the C/C++ code samples from the companion source code using the Windows 7 Driver Development Kit (DDK) build tools. The following steps detail this process. I strongly recommend that you complete these steps before you start reading, because you'll need the code samples to follow the experiments shown later in this book.

1. Download the Windows 7 DDK ISO image from the Microsoft Download Center at <http://www.microsoft.com/download/en/details.aspx?id=11800>, and save it to your local hard drive. Set plenty of time aside for this download if you have a slow Internet connection; the DDK ISO file is over 600 MB in size.



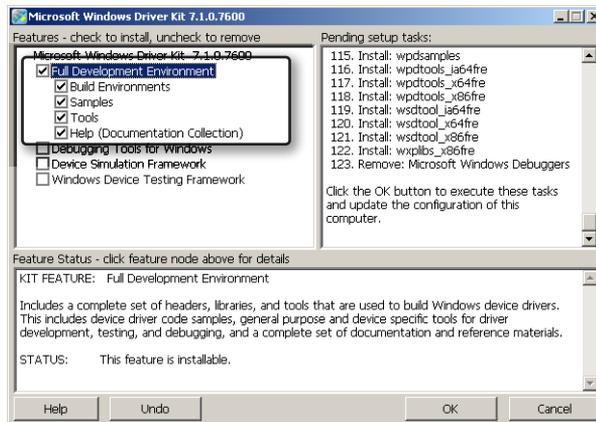
2. After the download is complete, mount the saved ISO file into a drive letter. There are several free tools for mounting ISO images on Windows. Virtual Clone Drive, which you can find on the Internet, is good freeware that works well both on Windows Vista and Windows 7. With that freeware installed, you should be able to right-click the ISO file and mount it, as shown in the following screen shot:



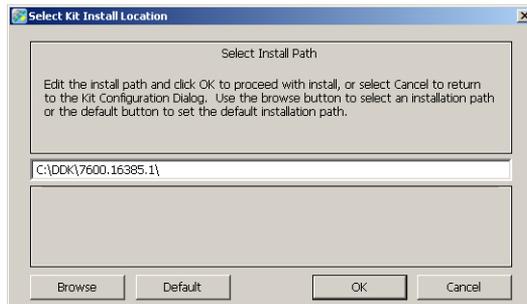
3. Double-click the newly mounted drive to kick off the DDK setup program, as shown here:



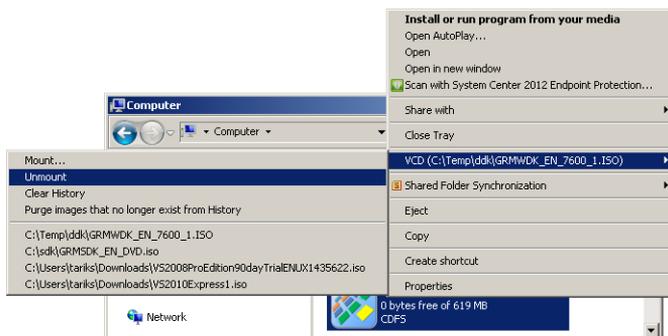
4. Select the Full Development Environment components from the DDK.



5. Then Install the components to the C:\DDK\7600.16385.1 directory, as shown in the following screen shot. This step will take several minutes to complete.



6. You can now unmount the DDK drive by right-clicking the mounted drive letter and selecting the Unmount menu action. This concludes this one-time installation of the Windows DDK build tools.



7. To build x86 binaries, first start a command prompt window and type the following command:

```
C:\DDK\7600.16385.1>bin\setenv.bat c:\DDK\7600.16385.1
```

8. You can then build all the native code samples at once. Simply navigate to the root directory that you extracted the companion source code to and issue the following command. It should take only a minute or so to build all the C/C++ code from the companion source code:

```
C:\book\code>bcz
```

Acknowledgments

I am indebted to Vance Morrison for reviewing my rough drafts. Vance has been a role model for me since I joined Microsoft, and his critical insight was tremendously helpful in improving the quality of my writing. I only hope that I came close to his high standards.

I'd also like to thank Silviu Calinoiu and Shawn Farkas for their detailed review of my draft chapters. Their attention to detail and precise feedback was tremendously helpful.

Kalin Toshev provided valuable technical feedback during the formative stages of this book, and his unwavering dedication to test-driven software development really inspired me to write this book. This book owes a lot to him.

Ajay Bhave, Cristian Levcovici, and Rico Mariani provided several ideas for improving the material covered in this book. This book certainly wouldn't be the same without their help and guidance.

Special thanks to Michael Fortin for writing the foreword, and to all my colleagues in the Windows fundamentals team for their support during the writing of this book.

Finally, I have to acknowledge my family, who play an important part in my life. I would like to dedicate this book to my parents for providing me with the opportunity to reach for my dreams and always being there when I needed them. Your love and support mean the world to me.

Errata & Book Support

We've made every effort to ensure the accuracy of this book and its companion content. Any errors that have been reported since this book was published are listed on our Microsoft Press site:

<http://www.microsoftpressstore.com/title/9780735662780>

If you find an error that is not already listed, you can report it to us through the same page.

If you need additional support, email Microsoft Press Book Support at *mspinput@microsoft.com*.

Please note that product support for Microsoft software is not offered through the addresses above.

We Want to Hear from You

At Microsoft Press, your satisfaction is our top priority, and your feedback our most valuable asset. Please tell us what you think of this book at:

<http://www.microsoft.com/learning/booksurvey>

The survey is short, and we read every one of your comments and ideas. Thanks in advance for your input!

Stay in Touch

Let's keep the conversation going! We're on Twitter: <http://twitter.com/MicrosoftPress>

Introduction

One exciting aspect of software programming is that there are usually many ways to accomplish the same goal. Unfortunately, this also presents software engineers with unique challenges when trying to make the best design or implementation choice for each situation. Experience plays a major role, and the learning process is often progressive as one learns to avoid past mistakes. Sadly, though, experience is often a variable concept. I have met several software engineers who, after spending a very long time working on an area, still lacked a basic understanding of how it *really* worked beyond the repetitive day-to-day tasks they grew accustomed to. I have also met others who have perfected their craft in a field after only a few years of working experience.

This book introduces a few techniques for methodically approaching software development problems primarily using the two “Swiss Army knives” of expert Microsoft Windows developers—namely, the Windows debuggers (WinDbg) and the Windows Performance Toolkit (Xperf). By focusing on *why* features and components work the way they do in the system rather than simply on *how* they work or *what* they do, this book tries to accelerate the process of learning by experience and to minimize the number of mistakes made when approaching new problems. An important part of the process is learning to compare and contrast with known solutions to existing ones.

Software engineering is still inherently a practical science. (Some might even argue it’s an art rather than a science.) While there is certainly no substitute for real experience, the topic can definitely be approached with the same methodical persistence that works so well for other scientific disciplines. In fact, this approach works even better in software engineering because all behaviors can be explained rationally. After all, it is all just code—whether it’s your own code or code written by others that you end up consuming in your software—and code can always be traced and understood.

Although this book deals with several architectural pillars of the Windows operating system as part of its debugging and tracing experiments, my main goal in writing it is less about covering those details and more about encouraging and developing this critical mindset. I hope to demonstrate how this approach can be used for solving a few interesting problems as part of this book, and that you continue to systematically apply debugging and tracing as you expand your learning beyond the topics directly covered here.

This book is not really about teaching native or managed code programming, either, although you’ll find several good coding examples in the companion source code. Because it takes an inside-out look at how to explore the system using debugging

and tracing tools, this book will probably appeal more to software engineers with the desire to understand system internals rather than those with a need to quickly learn how to make use of a specific technology. However, I believe the approach and mindset it aspires to inculcate are applicable regardless of the technology or level of expertise. In fact, contrary to what many think, the higher the level of technology involved, the harder it becomes to grasp what goes on behind the scenes and the more expertise is needed in order to investigate failures when things inevitably go awry and require debugging skills to save the day. In pure C, for example, a call to *malloc* is just that: a function call. In C++, a call to the *new* keyword is emitted by the compiler as a call to the *new* operator function to allocate memory for the object, followed by code to construct the object (again, emitted by the compiler to possibly initialize a virtual pointer and invoke the constructors of the base classes, construct member data objects, and finally invoke the user-provided constructor code for the target leaf class). In C# (.NET), things get even more involved, because a one-line call to the *new* keyword might involve compiling new code at runtime, performing security checks by the .NET execution engine, loading the modules where the target type is defined, tracking the object reference for later garbage collection, and so on.

Who Should Read This Book

This book is aimed at software engineers who desire to take their game to the next level, so to speak, and perfect their mastery of Windows as a development platform through the use of debugging and tracing tools.

Assumptions

Readers should have basic familiarity with the C/C++ and C# programming languages. A basic knowledge of the Win32 and .NET platforms is helpful, but not strictly required, because this book makes every effort to introduce basic concepts before expanding into more advanced topics.

Organization of This Book

This book is divided into three parts:

- Part 1, “A Bit of Background,” provides a brief overview of Windows development frameworks and the layers in the operating system that support them. This

basic knowledge is important to have when trying to make sense of the data surfaced by debugging and tracing tools.

- Part 2, “Debugging for Fun and Profit,” covers the architectural foundations of debuggers in the Windows operating system. It also presents a number of extensible strategies that will help you make the most of the Windows debuggers. In addition, this part also shows how to use the WinDbg debugger to better understand system internals by analyzing the important interactions between your code and the operating system.
- Part 3, “Observing and Analyzing Software Behavior,” continues this theme. It presents the Event Tracing for Windows (ETW) technology and illustrates how to leverage it in debugging and profiling investigations.
- Finally, you’ll find two short appendices at the end of the book that recap the most common debugging tasks and how to accomplish them using WinDbg.

The table of contents will help you locate chapters and sections quickly. In addition, each chapter starts with a list of the main points covered in the chapter, and concludes with a summary section. You can also download the source code for all the experiments and examples shown throughout the book.

Conventions in This Book

This book presents information using conventions (listed in the following table) designed to make the information readable and easy to follow:

Convention	Meaning
Sidebars	Boxed sidebars feature additional bits of information that might be helpful for a given subject.
Notes	Notes provide useful observations related to the content discussed in the main text.
Inline Code	Inline code—that is, code that appears within a paragraph—is shown in <i>italic</i> font.
Code Blocks	Code blocks are shown in a different font to help you distinguish code from text easily. Important statements appear in bold font to help you focus on those aspects of the code.
Debugger Listings	Debugger listings are shown in a different font, and important commands are bolded to highlight them. The listings are also often prefixed with the output from the standard <i>vertarget</i> debugger command, which displays the OS version and CPU architecture where the experiment was conducted.
Function Names	Function names are sometimes referenced using their WinDbg symbolic names. For example, <i>kernel32!CreateFileW</i> refers to the CreateFileW Win32 API (“W” for the Unicode flavor) exported by the <i>kernel32.dll</i> module.

System Requirements

You will need the following hardware and software to follow the experiments and code samples in this book:

- **Operating System** Windows Vista or later. Windows 7 (or Windows Server 2008 R2) is highly recommended.
- **Hardware** Any computer that supports the Windows 7 operating system (OS) requirements. Except for the live kernel debugging experiments, a second computer to serve as a host kernel-mode debugger machine is typically required for kernel debugging.



Note The target and host don't really need to be separate physical machines. A common kernel debugging configuration—detailed in Chapter 2, “Getting Started”—is to run the target machine where you conduct the experiments as a Windows 7 virtual OS on a Windows Server 2008 R2 physical host computer, and run the kernel debugger in the host OS.

- **Hard disk** 1 GB of free hard-disk space to download and save the Windows Software Development Kit (SDK) and Driver Development Kit (DDK) ISO images. 40 MB of free hard-disk space to download and compile the companion source code. An additional 3 GB is required for installing Microsoft Visual Studio 2010.
- **Software** The following tools are used in the debugging and tracing examples shown throughout this book:
 - Version 7.1 of the Windows 7 SDK, which can be downloaded from the Microsoft Download Center at <http://www.microsoft.com/download/en/details.aspx?id=8442>. Both the Windows Debuggers (WinDbg) and Windows Performance Toolkit (Xperf) are part of this SDK.
 - The Application Verifier tool, which can also be downloaded from the Microsoft Download Center at <http://www.microsoft.com/download/en/details.aspx?id=20028>.
 - The System Internals suite of developer tools, which can be downloaded from <http://technet.microsoft.com/en-us/sysinternals/bb842062>.
 - Visual Studio 2010, any edition (excluding the free, stripped-down Express edition). The Visual Studio Ultimate edition is preferred because some advanced features, such as static code analysis and performance profil-

ing, are not supported in other editions. The 90-day, free trial version of Visual Studio offered by Microsoft, which can be downloaded at <http://www.microsoft.com/download/en/details.aspx?id=12187>, should suffice.

- The Windows 7 Driver Development Kit (DDK) is used to compile the companion source code and can also be downloaded from the Microsoft Download Center at <http://www.microsoft.com/download/en/details.aspx?id=11800>.

Code Samples

Most of the chapters in this book include experiments and examples that let you interactively try out new material introduced in the main text. The programs used in these experiments can be downloaded from the following page:

<http://www.microsoftpressstore.com/title/9780735662780>

Follow the instructions to download the `Inside_Windows_Debugging_Samples.zip` file.

Installing the Code Samples

Follow these steps to install the code samples on your computer so that you can use them with the experiments and examples provided in this book:

1. Unzip the `Inside_Windows_Debugging_Samples.zip` file that you downloaded from the book's website into a folder named `\Book\Code`.



Warning Do not use directory names with spaces in the path hierarchy that you choose to download the samples to. The DDK build environment, which will be described shortly in the “Compiling the Code Samples” section, will fail to compile the source code if you do. It is recommended that you use `\Book\Code` as the root of the source code because this is the assumed location used when referencing the programs in the main text. The samples in the companion source code are organized by chapter, so there should be a folder for every chapter under this root path location.

2. If prompted, review the displayed end user license agreement. If you accept the terms, select the accept option, and then click Next.

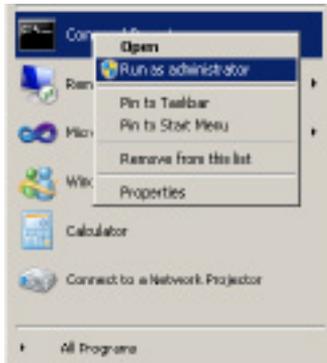


Note If the license agreement doesn't appear, you can access it from the same webpage from which you downloaded the `Inside_Windows_Debugging_Samples.zip` file.

Running the Code Samples

The code samples are organized by chapter and are referenced in the book main text by their respective directory path locations to help you find them easily.

Local Administrator rights are required by some sample programs. In Windows Vista and later, a command prompt must be launched as elevated in order to have full administrative privileges even when the user account is a member of the local built-in Administrators security group. To do so in Windows 7, for example, you need to right-click the Command Prompt menu item in the Windows Start menu, and select Run As Administrator, as shown in the following screen shot:



Compiling the Code Samples

The supporting programs used in the experiments presented in the book main text fall into three categories:

- **C++ samples** The binaries for these programs are deliberately omitted from the downloadable ZIP file. When you compile these code samples locally, WinDbg locates their symbols and source code files automatically. So the experiments presented in the early chapters of this book will “just work” in this configuration, without needing to specify the source and symbol locations explicitly in WinDbg. The steps to compile all the native C++ code samples at once are included later in this section.

- **C# (.NET) samples** For convenience, the compiled .NET programs are included in the downloadable ZIP file. You can use them as-is and you won't lose much, given WinDbg doesn't support source-level .NET debugging, but you can also follow the instructions included in the following section to recompile them if you prefer.
- **JavaScript and Visual Basic samples** These scripts are interpreted by the corresponding scripting engines and do not require compilation.

Compiling the .NET Code Samples

Compiling the .NET code samples from the companion source code requires version 4.0 of the Microsoft .NET Framework or later. Though this version isn't installed by default on Windows 7, .NET Framework 4.0 gets installed by many other dependent programs, such as Visual Studio 2010. You can also download and install a standalone version from the Microsoft download center at <http://www.microsoft.com/download/en/details.aspx?id=17851>.

Each C# code sample from the companion source code will have a helper compilation script under the same directory. This script uses the .NET 4.0 C# compiler directly and is easy to invoke, as illustrated in the following command:

```
C:\book\code\chapter_04\LoadException>compile.bat
```

If this script fails to find the C# compiler, you should verify that you've installed .NET 4.0 in the default directory location where this script expects it to exist. If .NET 4.0 exists but was installed to a different location on your system, you need to modify the script and provide that path instead.

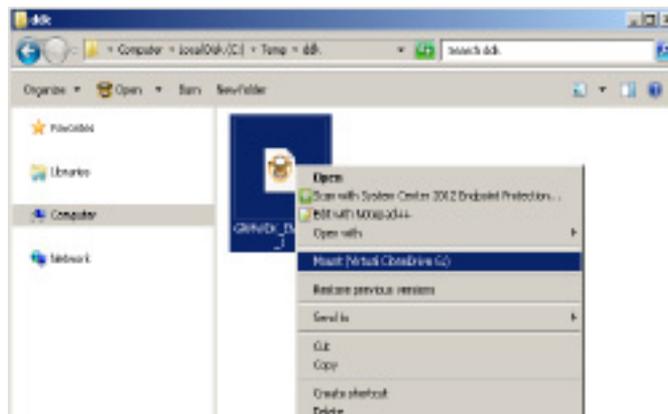
Compiling the C/C++ Code Samples

You can compile the C/C++ code samples from the companion source code using the Windows 7 Driver Development Kit (DDK) build tools. The following steps detail this process. I strongly recommend that you complete these steps before you start reading, because you'll need the code samples to follow the experiments shown later in this book.

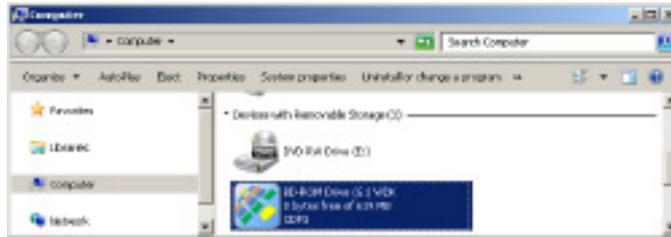
1. Download the Windows 7 DDK ISO image from the Microsoft Download Center at <http://www.microsoft.com/download/en/details.aspx?id=11800>, and save it to your local hard drive. Set plenty of time aside for this download if you have a slow Internet connection; the DDK ISO file is over 600 MB in size.



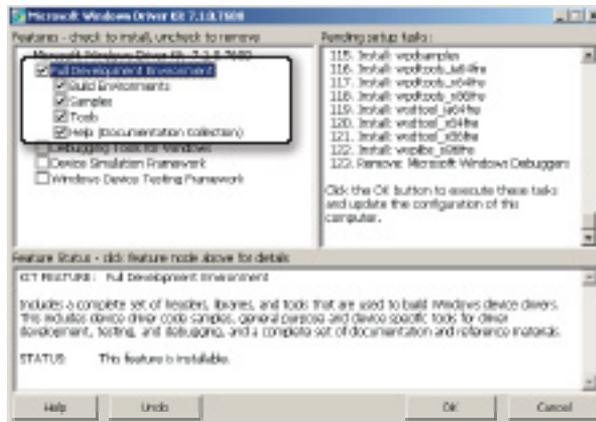
2. After the download is complete, mount the saved ISO file into a drive letter. There are several free tools for mounting ISO images on Windows. Virtual Clone Drive, which you can find on the Internet, is good freeware that works well both on Windows Vista and Windows 7. With that freeware installed, you should be able to right-click the ISO file and mount it, as shown in the following screen shot:



3. Double-click the newly mounted drive to kick off the DDK setup program, as shown here:



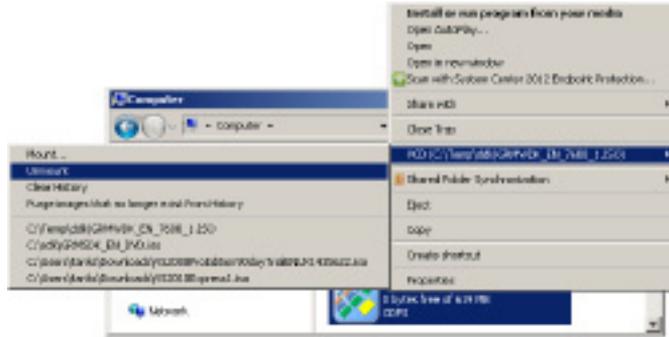
4. Select the Full Development Environment components from the DDK.



5. Then Install the components to the C:\DDK\7600.16385.1 directory, as shown in the following screen shot. This step will take several minutes to complete.



6. You can now unmount the DDK drive by right-clicking the mounted drive letter and selecting the Unmount menu action. This concludes this one-time installation of the Windows DDK build tools.



7. To build x86 binaries, first start a command prompt window and type the following command:

```
C:\DDK\7600.16385.1>bin\setenv.bat c:\DDK\7600.16385.1
```

8. You can then build all the native code samples at once. Simply navigate to the root directory that you extracted the companion source code to and issue the following command. It should take only a minute or so to build all the C/C++ code from the companion source code:

```
C:\book\code>bcz
```

Acknowledgments

I am indebted to Vance Morrison for reviewing my rough drafts. Vance has been a role model for me since I joined Microsoft, and his critical insight was tremendously helpful in improving the quality of my writing. I only hope that I came close to his high standards.

I'd also like to thank Silviu Calinoiu and Shawn Farkas for their detailed review of my draft chapters. Their attention to detail and precise feedback was tremendously helpful.

Kalin Toshev provided valuable technical feedback during the formative stages of this book, and his unwavering dedication to test-driven software development really inspired me to write this book. This book owes a lot to him.

Ajay Bhave, Cristian Levcovici, and Rico Mariani provided several ideas for improving the material covered in this book. This book certainly wouldn't be the same without their help and guidance.

Special thanks to Michael Fortin for writing the foreword, and to all my colleagues in the Windows fundamentals team for their support during the writing of this book.

Finally, I have to acknowledge my family, who play an important part in my life. I would like to dedicate this book to my parents for providing me with the opportunity to reach for my dreams and always being there when I needed them. Your love and support mean the world to me.

Errata & Book Support

We've made every effort to ensure the accuracy of this book and its companion content. Any errors that have been reported since this book was published are listed on our Microsoft Press site:

<http://www.microsoftpressstore.com/title/9780735662780>

If you find an error that is not already listed, you can report it to us through the same page.

If you need additional support, email Microsoft Press Book Support at *mspinput@microsoft.com*.

Please note that product support for Microsoft software is not offered through the addresses above.

We Want to Hear from You

At Microsoft Press, your satisfaction is our top priority, and your feedback our most valuable asset. Please tell us what you think of this book at:

<http://www.microsoft.com/learning/booksurvey>

The survey is short, and we read every one of your comments and ideas. Thanks in advance for your input!

Stay in Touch

Let's keep the conversation going! We're on Twitter: <http://twitter.com/MicrosoftPress>

PART 1

A Bit of Background

CHAPTER 1 Software Development in Windows 3

One of my esteemed mentors at Microsoft once told me the following story. He came back home one night to be greeted by a dismayed look on his wife's face. Her cherished wedding ring had just vanished down the bathroom sink drain and she was completely stumped. For all she knew, the wedding ring might very well have been in the ocean by the time they were having that conversation. To her, anything beyond the drain plug was a black box. My mentor, using the slight advantage of plumbing knowledge he had over his wife, knew about those nifty structures called "sink traps," which are usually J-shaped pipes located just beneath the sink. These are intended to "trap" a little bit of water and prevent sewer gas from flowing out of the drain pipes into the living space. But in addition, the traps also conveniently capture objects and keep them from going down the drain immediately. As it turned out, the ring was indeed in the sink trap, and he was able to retrieve it relatively easily.

The reason my mentor told me this story was to illustrate a close parallel in the world of software engineering: If you treat the APIs and frameworks you use in your code as pure black boxes, you might be able to get by for a little while, but you are certainly bound to experience some pretty anxious moments when you need to investigate failures that fall just outside of your own code, even if the solution—just like in the lost ring analogy—is right under your nose.

The first part of this book explores how your programs interact with the Microsoft Windows operating system (in a loose sense) and demonstrates why it's useful to have at least a cursory understanding of both those interactions and of the role of each subsystem. It also examines some important development frameworks shipped by Microsoft, and analyzes their positions relative to each other and to the operating system (OS) developer interfaces. It then concludes by introducing the Windows Software Development Kit (SDK) tools and, more specifically, the Swiss Army knives of expert Windows developers—namely, the Windows debugger (WinDbg) and the Windows Performance Toolkit (Xperf). This knowledge will serve as a perfect segue to the rest of the book, where you'll use debugging and tracing to write better software for the Windows operating system.

Software Development in Windows

In this chapter

Windows Evolution	3
Windows Architecture	7
Windows Developer Interface.	16
Microsoft Developer Tools.	28
Summary	30

Windows Evolution

Though this book focuses primarily on the post-Vista era of Windows, it's useful to look back at the history of Windows releases, because the roots of several building blocks of the underlying architecture can be traced all the way back to the Windows NT (an abbreviation for "New Technology") operating system. Windows NT was first designed and developed by Microsoft in the late '80s, and continued to evolve until its kernel finally became the core of all client and server versions of the Windows operating system.

Windows Release History

Windows XP marked a major milestone in the history of Windows releases by providing a unified code base for both the business (server) and consumer (client) releases of Windows. Though Windows XP was a client release (its server variant was Windows Server 2003), it technically succeeded both Windows 95/98/ME (a lineage of consumer operating systems that find their roots in the MS-DOS and Windows 3.1 operating systems) and Windows NT 4/Windows 2000, combining for the first time the power of the Windows NT operating system kernel and its robust architecture with many of the features that had made Windows 95 and Windows 98 instant hits with consumers and developers alike (friendly user design, aesthetic graphical interface, plug and play model, rich Win32 and DirectX API sets, and so on).

Though both the server and client releases of Windows now share the same kernel, they still differ in many of their features and components. (For example, only the server releases of Windows support multiple concurrent remote desktop user sessions.) Since the release of Windows XP in 2001, Windows Server has followed a release cycle that can be loosely mapped to corresponding Windows client releases. Windows Server 2003, for instance, shares many of the new kernel and API features that were added in Windows XP. Similarly, Windows Server 2008 R2 represents the server variant of Windows 7, which was released in late 2009. (Don't confuse this with Windows Server 2008, which is the server variant of Windows Vista.)

Figure 1-1 illustrates the evolution of the Windows family of operating systems, with their approximate release dates relative to each other.

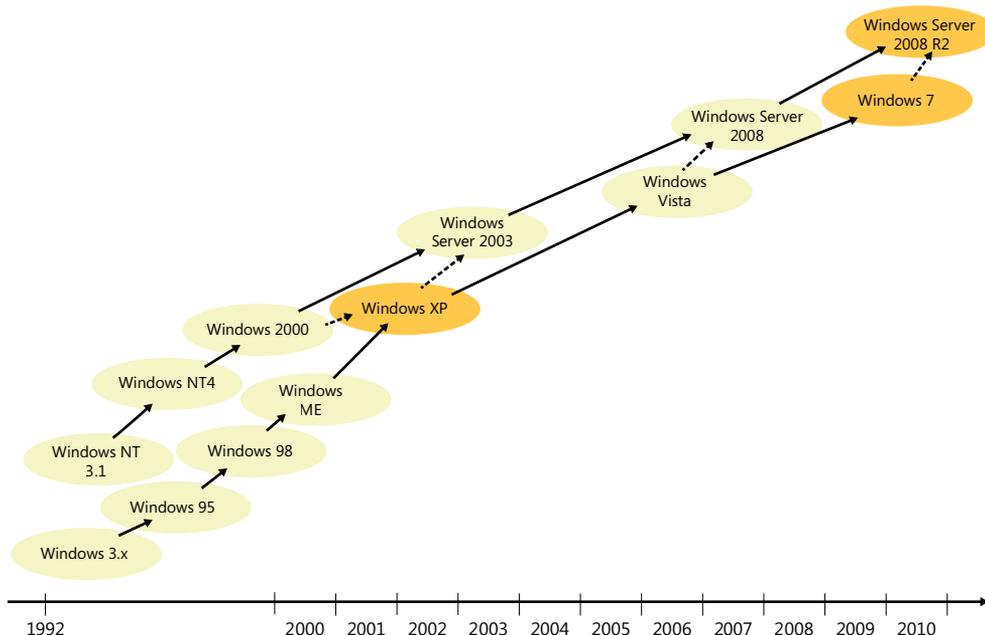


FIGURE 1-1 Timeline of major client and server releases of the Windows operating system since the early 90s.

Supported CPU Architectures

Windows was ported to many CPU architectures in the past. For example, Windows NT supported Alpha and MIPS processors until Windows NT 4. Windows NT 3.51 also had support for Power PC (another RISC family of processors that is used in many embedded devices, including, for example, Microsoft Xbox 360). However, Windows later narrowed its support to three CPU architectures: *x86* (a 32-bit family of processors, whose instruction set was designed by Intel), *x64* (also known as AMD-64, in reference to the fact this architecture was first introduced by AMD, though Intel also now releases processors implementing this instruction set), and *ia64* (another 64-bit instruction set designed by Intel in collaboration with Hewlett-Packard).

Microsoft shipped the first *ia64* version of Windows XP in late 2001 and followed it with an *x64* version in 2005. Microsoft later dropped support for *ia64* on client editions, including Windows XP. The *x86*, *x64*, and *ia64* architectures supported in Windows Server 2003 and Windows XP are exactly those that were also supported when Windows Server 2008 R2 and Windows 7 shipped at the end of 2009, though *x86* and *x64* are clearly the more widely used Windows architectures nowadays. Note, however, that Windows Server no longer supports *x86*; it now supports only 64-bit architectures. Also, Microsoft announced early in 2011 that its upcoming release of the Windows operating system will be capable of running on *ARM* (in addition to the *x86* and *x64* platforms), a RISC instruction set that's widely used in embedded utilities, smartphones, and tablet (slate) devices thanks in large part to its efficient use of battery power.

Understanding the underlying CPU architecture of the Windows installation you are working on is very important during debugging and tracing because you often need to use native tools that correspond to your CPU architecture. In addition, sometimes you will also need to understand the disassembly of the code you are analyzing in the debugger, which is different for each CPU. This is one reason many debugger listings in this book also show the underlying CPU architecture that they were captured on so that you can easily conduct any further disassembly inspection you decide to do on the right target platform. In the following listing, for example, the *vertarget* command shows a Windows 7 AMD64 (*x64*) operating system. You'll see more about this command and others in the next chapter, so don't worry about how to issue it for now.

```
lkd> $ Multi-processor (2 processors or cores) Windows 7 x64 system
lkd> vertarget
Windows 7 Kernel Version 7600 MP (2 procs) Free x64
Built by: 7600.16385.amd64fre.win7_rtm.090713-1255
```

Given the widespread use of *x86* and *x64*, and because you can also execute *x86* programs on *x64* machines, the majority of experiments in this book are conducted using the *x86* architecture so that you can follow them the way they're described in this book regardless of your target architecture. Though *x86* has been the constant platform of choice for Windows since its early days in the '80s, 64-bit processors continue to gain in popularity even among home computers and laptops, which now often carry *x64* versions of Windows 7.

Windows Build Flavors

The *vertarget* debugger command output shown in the previous section referred to the Windows version on the target machine as a "free" (also known as *retail*) build. This flavor is the only one ever shipped to end users by Microsoft for any of the supported processor architectures. There is, however, another flavor called a "checked" (also known as *debug*) build, which MSDN subscribers can obtain from Microsoft if they want to test the software they build with this flavor of the Windows operating system. It's important to realize that checked flavors are mostly meant to help driver developers; they don't derive their name at all from being "tested"—or otherwise "checked"—more thoroughly than the free flavors.

If you recall, the Introduction of this book recommended using the Driver Development Kit (DDK) build environment if you wanted to recompile the companion C++ sample code. As was explained then, you can also specify the build flavor you want to target (the default being *x86 free* in the Windows 7 DDK) when starting a DDK build environment. This is, in fact, also how the checked flavor of Windows is built internally at Microsoft because the same build environment made available in the DDK is also used by Windows developers to compile the Windows source code. For example, the following command starts a DDK build environment where your source code is compiled into *x64* binaries using the checked (*chk*) build flavor. This essentially turns off a few compiler optimizations and defines debug build macros (such as the `DBG` preprocessor variable) that turn on “debug” sections of the code, including assertions (such as the `NT_ASSERT` macro).

```
C:\DDK\7600.16385.1>bin\setenv.bat c:\DDK\7600.16385.1 chk x64
```

Naturally, you don’t really need a checked build of Windows to run your checked binaries, and the main difference between your “free” and “checked” binaries is that the assertions you put in your code will occur only in the “checked” flavor. The benefit of the checked flavor of Windows itself is that it also contains many additional assertions in the system code that can point out implementation problems in your code, which is usually useful if you are developing a driver. The drawback to that Windows flavor, of course, is that it runs much slower than the free flavor and also that you must run it with a kernel debugger attached at all times so that you can ignore assertions when they’re hit; otherwise, those assertions might go unhandled and cause the machine to crash and reboot if they are raised in code that runs in kernel mode.

Windows Servicing Terminology

Each major Windows release is usually preceded by a few public milestones that provide customers with a preview of the features included in that release. Those prerelease milestones are usually called *Alpha*, *Beta1*, *Beta2*, and *RC* or release candidate, in this chronological order, though several Windows releases have either skipped some of these milestones or named them differently. These prerelease milestones also present an opportunity for Microsoft to engage with customers and collect their feedback before Windows is officially “released to manufacturing,” a milestone referred to as *RTM*.

You will again recognize the major version of Windows in the build information displayed by the `vertarget` command that accompanies many of the debugger listings presented in this book. For example, the following listing shows that the target machine is running Windows 7 RTM and that July 13, 2009 (identified by the “090713” substring in the following output) is the date this particular Windows build was produced at Microsoft.

```
Tkd> vertarget
Windows 7 Kernel Version 7600 MP (2 procs) Free x64
Built by: 7600.16385.amd64fre.win7_rtm.090713-1255
```

In addition to the major client and server releases for each Windows operating system, Microsoft also ships several servicing updates in between those releases that get automatically delivered via the Windows Update pipeline and usually come in one of the following forms:

- **Service packs** These releases usually occur a few years apart after RTM and compile the smaller updates made in between (and, on occasion, new features requested by customers) into a single package that can be applied at once by both consumers and businesses. They are often referred to using the “SP” abbreviation, followed by the service pack number. For example, SP1 is the first service pack after RTM, SP2 is the second, and so on.

Service packs are considered major releases and are subjected to the same rigorous release process that accompanies RTM releases. In fact, many Windows service packs also have one or more release candidate (RC) milestones before they’re officially released to the public. The number of service packs for a major Windows release is often determined by customer demand and also the amount of changes accumulated since the last service pack. Windows NT4, for example, had six service packs, while Windows Vista had two.

- **GDR updates** GDR (General Distribution Release) updates are issued to address bugs with broad impact or security implications. The frequency varies by need, but these updates are usually released every few weeks. These fixes are also rolled up into the following service pack release.

For example, the following output indicates that the target debugging machine is running a version of Windows 7 SP1. Notice also that the version of *kernel32.dll* that’s installed on this machine comes from a GDR update subsequent to the initial Windows 7 SP1 release.

```
0:000> vertarget
Windows 7 Version 7601 (Service Pack 1) MP (2 procs) Free x86 compatible
kernel32.dll version: 6.1.7601.17651 (win7sp1_gdr.110715-1504)
```

Windows Architecture

The fundamental design of the Windows operating system, with an executive that runs in kernel mode and a complementary set of user-mode system support processes (*smss.exe*, *csrss.exe*, *winlogon.exe*, and so on) to help manage additional system facilities, has for the most part remained unchanged since the inception of the Windows NT operating system back in the late '80s. Each new version of Windows naturally brings about a number of new components and APIs, but understanding how they fit in the architectural stack often starts with knowing how they interact with these core components of the operating system.

Kernel Mode vs. User Mode

Kernel mode is an execution mode in the processor that grants access to all system memory (including user-mode memory) and unrestricted use of all CPU instructions. This CPU mode is what enables the Windows operating system to prevent *user-mode* applications from causing system instability by accessing protected memory or I/O ports.

Application software usually runs in user mode and is allowed to execute code in kernel mode only via a controlled mechanism called a *system call*. When the application wants to call a system service exposed by code in the OS that runs in kernel mode, it issues a special CPU instruction to switch the calling thread to kernel mode. When the service call completes its execution in kernel mode, the operating system switches the thread context back to user mode, and the calling application is able to continue its execution in user mode.

Third-party vendors can get their code to run directly in kernel mode by implementing and installing signed *drivers*. Note that Windows is a monolithic system in the sense that the OS kernel and drivers share the same address space, so any code executing in kernel mode gets the same unrestricted access to memory and hardware that the core of the Windows operating system would have. In fact, several parts of the operating system (the NT file system, the TCP/IP networking stack, and so on) are also implemented as drivers rather than being provided by the kernel binary itself.

The Windows operating system uses the following layering structure for its kernel-mode operations:

- **Kernel** Implements core low-level OS services such as thread scheduling, multiprocessor synchronization, and interrupt/exception dispatching. The kernel also contains a set of routines that are used by the executive to expose higher-level semantics to user-mode applications.
- **Executive** Also hosted by the same “kernel” module in Windows (NTOSKRNL), and performs base services such as process/thread management and I/O dispatching. The executive exposes documented functions that can be called from kernel-mode components (such as drivers). It also exposes functions that are callable from user mode, known as *system services*. The typical entry point to these executive system services in user mode is the *ntdll.dll* module. (This is the module that has the system call CPU instruction!) During these system service calls, the executive allows user-mode processes to reference the objects (process, thread, event, and so on) it implements via indirect abstractions called object *handles*, which the executive keeps track of using a per-process handle table.
- **Hardware Abstraction Layer** The HAL (*hal.dll*) is a loadable kernel-mode module that isolates the kernel, executive, and drivers from hardware-specific differences. This layer sits at the very bottom of kernel layers and handles key hardware differences so that higher-level components (such as third-party device drivers) can be written in a platform-agnostic way.
- **Windows and Graphics Subsystem** The Win32 UI and graphics services are implemented by an extension to the kernel (*win32k.sys* module) and expose system services for UI applications. The typical entry point to these services in user mode is the *user32.dll* module.

Figure 1-2 illustrates this high-level architecture.

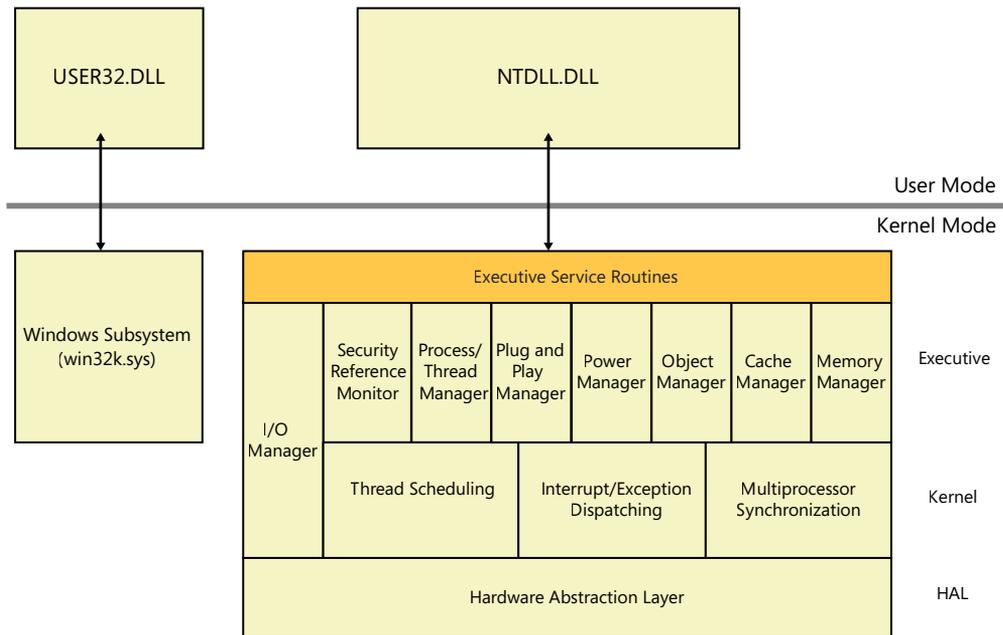


FIGURE 1-2 Kernel-mode layers and services in the Windows operating system.

User-Mode System Processes

Several core facilities (logon, logoff, user authentication, and so on) of the Windows operating system are primarily implemented in user mode rather than in kernel mode. A fixed set of user-mode system processes exists to complement the OS functionality exposed from kernel mode. Here are a few important processes that fall in this category:

- ***Smss.exe*** User sessions in Windows represent resource and security boundaries and offer a virtualized view of the keyboard, mouse, and physical display to support concurrent user logons on the same OS. The state that backs these sessions is tracked in a kernel-mode virtual memory space usually referred to as the *session space*. In user mode, the session manager subsystem process (*smss.exe*) is used to start and manage these user sessions.

A “leader” *smss.exe* instance that’s not associated with any sessions gets created as part of the Windows boot process. This leader *smss.exe* creates a transient copy of itself for each new session, which then starts the *winlogon.exe* and *csrss.exe* instances corresponding to that user session. Although having the leader session manager use copies of itself to initialize new sessions doesn’t provide any practical advantages on client systems, having multiple *smss.exe* copies running concurrently can provide faster logon of multiple users on Windows Server systems acting as Terminal Servers.

- ***Winlogon.exe*** The Windows logon process is responsible for managing user logon and logoff. In particular, this process starts the logon UI process that displays the logon screen when the user presses the Ctrl+Alt+Del keyboard combination and also creates the processes

responsible for displaying the familiar Windows desktop after the user is authenticated. Each session has its own instance of the *winlogon.exe* process.

- **Csrss.exe** The client/server runtime subsystem process is responsible for the user-mode portion of the Win32 subsystem (*win32k.sys* being the kernel-mode portion) and also was used to host the UI message loop of console applications prior to Windows 7. Each user session has its own instance of this process.
- **Lsass.exe** The local security authority subsystem process is used by *winlogon.exe* to authenticate user accounts during the logon sequence. After successful authentication, LSASS generates a security access token object representing the user's security rights, which are then used to create the new explorer process for the user session. New child processes created from that shell then inherit their access tokens from the initial explorer process security token. There is only one single instance of this process, which runs in the noninteractive session (known as session 0).
- **Services.exe** This system process is called the *NT service control manager* (SCM for short) and runs in session 0 (noninteractive session). It's responsible for starting a special category of user-mode processes called *Windows services*. These processes are generally used by the OS or third-party applications to carry out background tasks that do not require user interaction. Examples of Windows services include the spooler print service (*spooler*); the task scheduler service (*schedule*); the COM activation services, also known as the COM SCM (*RpcSs* and *DComLaunch*); and the Windows time service (*w32time*).

These processes can choose to run with the highest level of user-mode privileges in Windows (*LocalSystem* account), so they are often used to perform privileged tasks on behalf of user-mode applications. Also, because these special processes are always started and stopped by the SCM process, they can be started on demand and are guaranteed to have at most one active instance running at any time.

All of the aforementioned system-support processes run under the *LocalSystem* account, which is the highest privileged account in Windows. Processes that run with this special account identity are said to be a part of the trusted computing base (TCB) because once user code is able to run with that level of privilege, it is also able to bypass any checks by the security subsystem in the OS.

User-Mode Application Processes

Every user-mode process (except for the leader *smss.exe* process mentioned earlier) is associated with a user session. These user-mode processes are boundaries for a memory address space. As far as scheduling in Windows is concerned, however, the most fundamental scheduling units remain the threads of execution and processes are merely containers for those threads. It's also important to realize that user-mode processes (more specifically, the threads they host) also often run plenty of code in kernel mode. Although your application code might indeed run in user mode, it's often the case that it also calls into system services (through API layers that call down to NTDLL or USER32 for the system call transitions) that end up transitioning to kernel mode on your behalf. This is why it makes sense to always think of your software (whether it's user-mode software or kernel drivers) as

an extension of the Windows operating system and also that you understand how it interacts with the “services” provided by the OS.

Processes, in turn, can be placed in containers called *job objects*. These executive objects can be very useful to manage a group of processes as a single unit. Unlike threads and processes, job objects are often overlooked when studying the Windows architecture despite their unique advantages and the useful semantics they provide. Figure 1-3 illustrates the relationship between these fundamental objects.

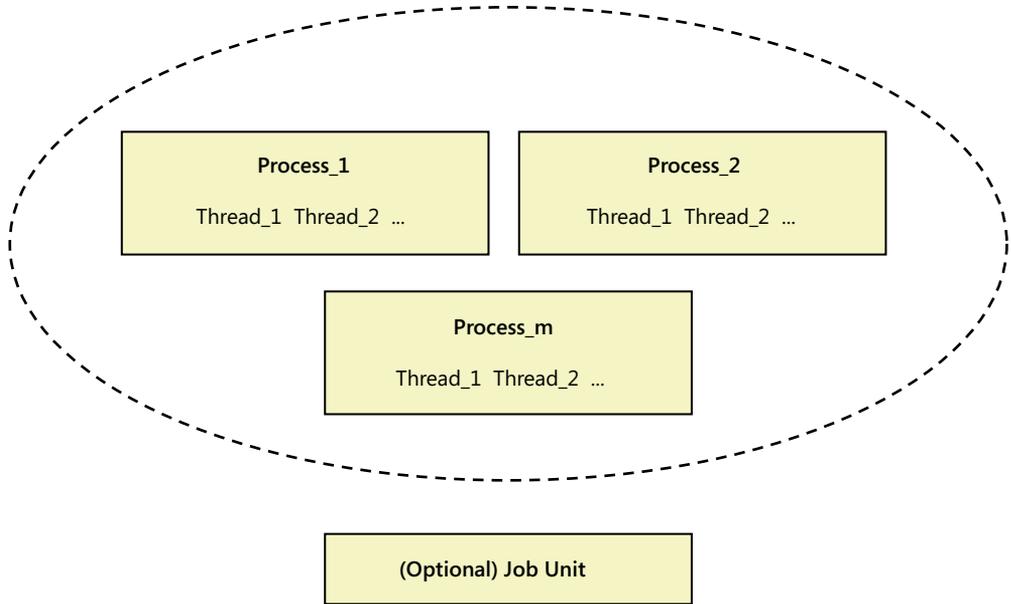


FIGURE 1-3 Threads, processes, and jobs in Windows.

Job objects can be used to provide common execution settings for a set of processes and, among other things, to control the resources used by member processes (such as the amount of memory consumed by the job and the processors used for its execution) or their UI capabilities.

One particularly useful feature of job objects is that they can be configured to terminate their processes when their user-mode job handle is closed (either using an explicit *kernel32!CloseHandle* API call, or implicitly when the kernel runs down the handles in the process handle table when the process kernel object is destroyed). To provide a practical illustration, the following C++ program shows how to take advantage of the job-object construct exposed by the Windows executive in a C++ user-mode application to start a child (“worker”) process and synchronize its lifetime with that of its parent process. This is often useful in the case of worker processes whose sole purpose is to serve requests in the context of their parent process, in which case it becomes critical not to “leak” those worker instances should the parent process die unexpectedly. (The reverse is more straightforward because the parent process can easily monitor when the child dies by simply waiting on the worker process handle to become signaled using the *kernel32!WaitForSingleObject* Win32 API.)

To follow this experiment, remember to refer back to the Introduction of this book, which contains step-by-step instructions for how to build the companion source code.

```
//
// C:\book\code\chapter_01\WorkerProcess>main.cpp
//
class CMainApp
{
public:
    static
    HRESULT
    MainHR()
    {
        HANDLE hProcess, hPrimaryThread;
        CHandle shProcess, shPrimaryThread;
        CHandle shWorkerJob;
        DWORD dwExitCode;
        JOB_OBJECT_EXTENDED_LIMIT_INFORMATION exLimitInfo = {0};
        CStringW shCommandLine = L"notepad.exe";

        ChkProlog();

        //
        // Create the job object, set its processes to terminate on
        // handle close (similar to an explicit call to TerminateJobObject),
        // and then add the current process to the job.
        //
        shWorkerJob.Attach(CreateJobObject(NULL, NULL));
        ChkWin32(shWorkerJob);

        exLimitInfo.BasicLimitInformation.LimitFlags =
            JOB_OBJECT_LIMIT_KILL_ON_JOB_CLOSE;
        ChkWin32(SetInformationJobObject(
            shWorkerJob,
            JobObjectExtendedLimitInformation,
            &exLimitInfo,
            sizeof(exLimitInfo)));

        ChkWin32(AssignProcessToJobObject(
            shWorkerJob,
            ::GetCurrentProcess()));

        //
        // Now launch the new child process (job membership is inherited by default)
        //
        wprintf(L"Launching child process (notepad.exe) ...\n");
        ChkHr(LaunchProcess(
            shCommandLine.GetBuffer(),
            0,
            &hProcess,
            &hPrimaryThread));
    }
};
```

```

shProcess.Attach(hProcess);
shPrimaryThread.Attach(hPrimaryThread);

//
// Wait for the worker process to exit
//
switch (WaitForSingleObject(shProcess, INFINITE))
{
    case WAIT_OBJECT_0:
        ChkWin32(::GetExitCodeProcess(shProcess, &dwExitCode));
        wprintf(L"Child process exited with exit code %d.\n", dwExitCode);
        break;
    default:
        ChkReturn(E_FAIL);
}

    ChkNoCleanup();
}
};

```

One key observation here is that the parent process is assigned to the new job object before the new child process is created, which allows the worker process to automatically inherit this job membership. This means in particular that there is no time window in which the new process would exist without being a part of the job object. If you kill the parent process (using the Ctrl+C signal, for example), you will notice that the worker process (*notepad.exe* in this case) is also terminated at the same time, which was precisely the desired behavior.

```

C:\book\code\chapter_01\WorkerProcess>objfre_win7_x86\i386\workerprocess.exe
Launching child process (notepad.exe) ...
^C

```

Low-Level Windows Communication Mechanisms

With code executing in kernel and user modes, and also inside the boundaries of per-process address spaces in user mode, the Windows operating system supports several mechanisms for allowing components to communicate with each other.

Calling Kernel-Mode Code from User Mode

The most basic way to call kernel-mode code from user-mode components is the *system call* mechanism mentioned earlier in this chapter. This mechanism relies on native support in the CPU to implement the transition in a controlled and secure manner.

One inherent drawback to the system call mechanism is that it relies on a hard-coded table of well-known executive service routines to dispatch the request from the client code in user mode to its intended target service routine in kernel mode. This doesn't extend well to kernel extensions implemented in the form of drivers, however. For those cases, another mechanism—called *I/O control*

commands (IOCTL)—is supported by Windows to enable user-mode code to communicate with kernel-mode drivers. This is done through the generic *kernel32!DeviceIoControl* API, which takes the user-defined IOCTL identifier as one of its parameters and also a handle to the device object to which to dispatch the request. The transition to kernel mode is still performed in the NTDLL layer (*ntdll!NtDeviceIoControlFile*) and internally also uses the system call mechanism. So, you can think of the IOCTL method as a higher-level user/kernel communication protocol built on top of the raw system call services provided by the OS and CPU.

Internally, I/O control commands are processed by the I/O manager component of the Windows executive, which builds what is called an *I/O request packet (IRP for short)* that it then routes to the device object requested by the caller from user mode. IRP processing in the Windows executive uses a layered model where devices have an associated driver stack that handles their requests. When an IRP is sent to a top-level device object, it travels through its device stack starting at the top, passing through each driver in the corresponding device stack and giving it a chance to either process or ignore the command. In fact, IRPs are also used in kernel mode to send commands to other drivers so that the same IRP model is used for interdriver communication in the kernel. Figure 1-4 depicts this architecture.

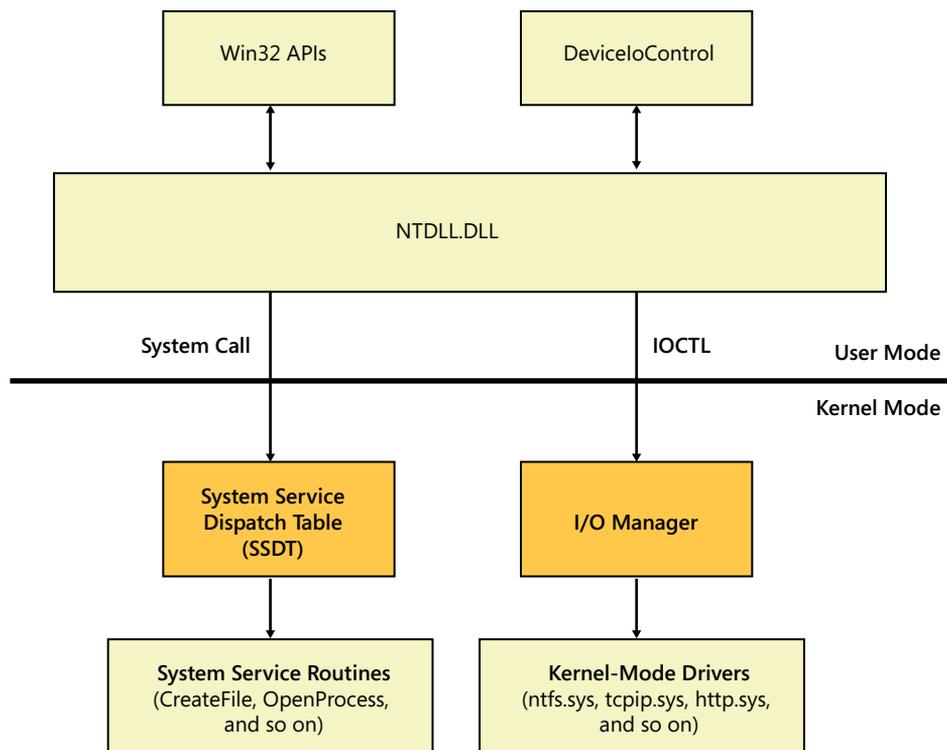


FIGURE 1-4 User-mode to kernel-mode communication mechanisms.

Calling User-Mode Code from Kernel Mode

Code that runs in kernel mode has unrestricted access to the entire virtual address space (both the user and kernel portions), so kernel mode in theory could invoke any code running in user mode. However, doing so requires first picking a thread to run the code in, transitioning the CPU mode back to user mode, and setting up the user-mode context of the thread to reflect the call parameters. Fortunately, however, only the system code written by Microsoft really needs to communicate with random threads in user mode. The drivers you write, on the other hand, need to call back to user mode only in the context of a device IOCTL initiated by a user-mode thread, so they do not need a more generic kernel-mode to user-mode communication mechanism.

A standard way for the system to execute code in the context of a given user-mode thread is to send an asynchronous procedure call (APC) to that thread. For example, this is exactly how thread suspension works in Windows: the kernel simply sends an APC to the target thread and asks it to execute a function to wait on its internal thread semaphore object, causing it to become suspended. APCs are also used by the system in many other scenarios, such as in I/O completion and thread pool callback routines, just to cite a couple.

Interprocess Communication

Another way for communicating between user-mode processes and code in kernel mode, as well as between user-mode processes themselves, is to use the advanced local procedure call (ALPC) mechanism. ALPC was introduced in the Windows Vista timeframe and is a big revision of the LPC mechanism, a feature that provided in many ways the bloodline of low-level intercomponent communication in Windows since its early releases.

ALPC is based on a simple idea: a server process first opens a kernel port object to receive messages. Clients can then connect to the port if allowed by the server owning the port and start sending messages to the server. They are also able to wait until the server has fetched and processed the message from the internal queue that's associated with the ALPC port object.

In the case of user/user ALPC, this provides a basic low-level interprocess communication channel. In the case of kernel/user ALPC channels, this essentially provides another (indirect) way for user-mode applications to call code in kernel mode (whether it's in a driver or in the kernel module itself) and vice versa. An example of this communication is the channel that's established between the *lsass.exe* user-mode system process and the security reference monitor (SRM) executive component in kernel mode, which is used, for example, to send audit messages from the executive to *lsass.exe*. Figure 1-5 illustrates this architecture.

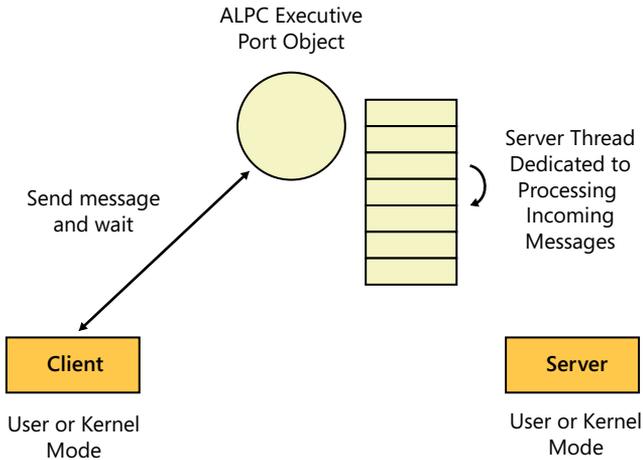


FIGURE 1-5 ALPC communication in Windows.

ALPC-style communication is used extensively in the operating system itself, most notably as it pertains to this book to implement the low-level communication protocol that native user-mode debuggers employ to receive various debug events from the process they debug. ALPC is also used as a building block in higher-level communication protocols such as local RPC, which in turn is used as the transport protocol in the COM model to implement interprocess method invocations with proper parameter marshaling.

Windows Developer Interface

Developers can extend the services that ship with the Windows operating system by building extensions in the way of kernel drivers or standalone user-mode Windows applications. This section examines some of the key layers and APIs that make building these extensions possible.

Developer Documentation Resources

Microsoft documents several APIs that developers can use when building their applications. The difference between these published interfaces and the internal (private) implementation details of the platform is that Microsoft has committed over the past two decades to building an ecosystem in which the public interfaces are carried forward with new releases of Windows, affording application developers the confidence that the applications they build today will continue to work on future OS versions. It's fair to say that this engineering discipline is one of the reasons Windows has been so successful with developers and end users alike.

Microsoft documents all of the interfaces and APIs it publishes on the Microsoft Developer Network (MSDN) website at <http://www.microsoft.com/msdn>. When writing your software, you should use only officially supported public APIs so that your software doesn't break when new versions of the operating system are released by Microsoft. Undocumented APIs can often disappear or get

renamed, even in service pack releases of the same OS version, so you should never target them in your software unless you are prepared to deal with the wrath of your customers when your software ominously stops working after a new Windows update is installed. That's not to say you shouldn't be interested in internal implementation details, of course. In fact, this book proves that knowing some of those details is often important when debugging your programs and can help you analyze their behaviors more proficiently, which in turns also helps you write better and more reliable software in the long run.

In addition to documenting the public APIs written by Microsoft for developers, the MSDN website also contains a wealth of articles describing features or areas at a higher level. In particular, it hosts a special category of articles, called *Knowledge Base articles* (KB articles for short) that are published by Microsoft's customer support team to document workarounds for known issues. Generally speaking, the MSDN website should be your first stop when looking up the documented behavior of the APIs you use in your code or when trying to learn how to use a new feature in the OS.

WDM, KMDF, and UMDF

Developers can run their code with kernel-mode privileges and extend the functionality of the OS by implementing a kernel-mode driver. Though the vast majority of developers will never need to write a kernel driver, it's still useful to understand the layered plug-in models used by Windows to support these kernel extensions because this knowledge sometimes can help you make sense of kernel call stacks during kernel-mode debugging investigations.

Driver extensions are often needed to handle communication with hardware devices that aren't already supported because—as mentioned earlier—user-mode code isn't allowed to access I/O ports directly. In addition, drivers are sometimes used to implement or extend system features. For example, many tools in the SysInternals suite—including the process monitor tool, which installs a filter driver to monitor access to system resources—internally install drivers when they're executed to implement their functionality.

There are many ways to write drivers, but the main model used to implement them in Windows is the Windows Driver Model (WDM). Because this model asks a lot from driver developers in terms of handling all the interactions with the I/O manager and the rest of the operating system and often results in a lot of duplicated boilerplate code that has to be implemented by all driver developers, the kernel-mode driver framework (KMDF) was introduced to simplify the task of writing kernel-mode drivers. Keep in mind, however, that KMDF doesn't replace WDM; rather, it's a framework that helps you more easily write drivers that comply with WDM's requirements. Generally speaking, you should write your drivers using KMDF unless you find a good reason not to do so, such as when you need to write non-WDM drivers. This is the case, for instance, for network, SCSI, or video drivers, which have their own world, so to speak, and require you to write what is called a "miniport" driver to plug into their respective port drivers.

A subset of hardware drivers also can be executed completely in user mode (though without direct access to kernel memory space or I/O ports). These drivers can be developed using another framework shipped by Microsoft called the user-mode driver framework, or UMDF. For more details on

the different driver models and their architectures, you can find a wealth of resources on the MSDN website at <http://msdn.microsoft.com>. The OSR website at <http://www.osronline.com> is also worth a visit if you ever need to write or debug drivers in Windows.

The NTDLL and USER32 Layers

As mentioned earlier in this chapter, the NTDLL and USER32 layers contain the entry points to the executive service routines and kernel-mode portion of the Win32 subsystem (*win32k.sys*), respectively.

There are hundreds of executive service stubs in the NTDLL module (*ntdll!NtSetEvent*, *ntdll!NtReadFile*, and many others). The majority of these service stubs are undocumented in the MSDN, but a few stub entry points were deemed generally useful to third-party system software and are documented. The NTDLL.DLL system module also hosts several low-level OS features, such as the module loader (*ntdll!Ldr** routines), the Win32 subsystem process communication functions (*ntdll!Csr** routines), and several run-time library functions (*ntdll!Rtl** routines) that expose features such as the Windows heap manager and Win32 critical section implementations.

The NTDLL module is used by many other DLLs in the Win32 API to transition into kernel mode and call executive service routines. Similarly, the USER32 DLL is also used by the Windows graphics architectural stack (DirectX, GDI32, and so on) as the gateway to transition into kernel mode so that it can communicate with the graphics processing unit (GPU) hardware.

The Win32 API Layer

The Win32 API layer is probably the most important layer to learn for developers who are new to Windows because it's the official public interface to the services exposed by the operating system. All of the Win32 API functions are documented in the MSDN, along with their expected parameters and potential return codes. Even if you are writing your software using a higher-level development framework or API set, as most of us do these days, being aware of the capabilities exposed at this layer will help you get a much better feel for a framework's advantages and limitations relative to other choices, as well as the raw capabilities exposed by the Win32 API and Windows executive.

The Win32 API layer covers a large set of functionality, going from basic services like creating threads/processes or drawing shapes on the screen to higher-level areas such as cryptography. The most basic services at the bottom of the Win32 API's architectural stack are exposed in the *kernel32.dll* module. Other widely used Win32 DLL modules are *advapi32.dll* (general utility functions), *user32.dll* (Windows and user object functions), and *gdi32.dll* (graphics functions). In Windows 7, the Win32 DLL modules are now layered so that lower-level base functions aren't allowed to call up to higher-level modules in the hierarchical stack. This layering engineering discipline helps prevent circular dependencies between modules and also minimizes the performance impact of bringing a new DLL dependency from the Win32 API set into your process address space. This is why you will see that many of the public APIs exported in the *kernel32.dll* module now simply forward their calls to the implementation defined in the lower-level *kernelbase.dll* DLL module, which is useful to know when trying to set system breakpoints in the debugger. This layered hierarchy is demonstrated in Figure 1-6.

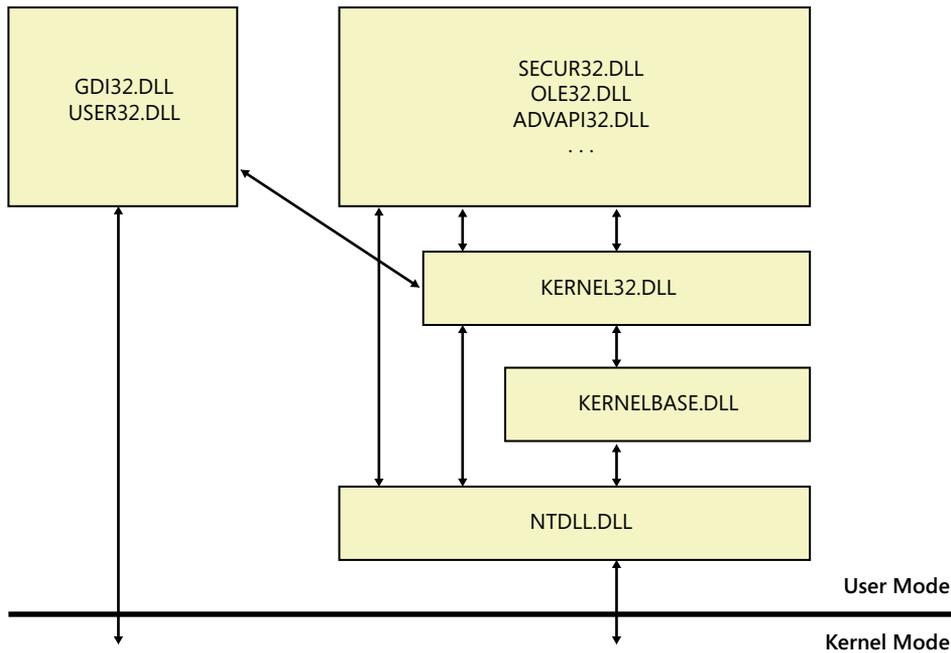


FIGURE 1-6 Low-level Win32 DLL modules.

The COM Layer

The Component Object Model (COM) was introduced by Microsoft in the mid-90s as a user-mode framework to enable developers to write reusable object-oriented components in different programming languages. If you are new to COM, the best resource to use to get started and gain an understanding of its model is the “Component Object Model Specification” document, which you can still find online. Though this is an old (circa 1995) and relatively long document, it provides a great overview of the COM binary specification, and much of what it describes still holds true to this day.

Over time, the use of the term “COM” grew to cover a number of different but related technologies, including the following:

- The object model itself, which the label “COM” was technically designed to describe at first. Key parts of this object model are the standard *IUnknown* and *IClassFactory* interfaces, the idea of separation of class interfaces (the public contract) from the internal COM class implementation, and the ability to query a server object for an implementation of the contracts that it supports (*IUnknown::QueryInterface*). This model, in its pure form, is one of the most elegant contributions by Microsoft to the developer ecosystem, and it has had a deep impact that still reverberates to this day in the form of various derivative technologies based on that model.
- The interprocess communication protocol and registration that allows components to communicate with each other without the client application having to know where the server component lives. This enables COM clients to transparently use servers implemented

locally in-process (DLLs) or out-of-process (EXEs), as well as servers that are hosted on a remote computer. The distinction between COM and Distributed COM (DCOM, or COM across machines) is often only theoretical, and most of the internal building blocks are shared between the two technologies.

- The protocol specifications built on top of the COM object model to enable hosts to communicate with objects written in multiple languages. This includes the OLE Automation and ActiveX technologies, in particular.

COM in the Windows Operating System

COM is omnipresent in the Windows operating system. Although the Microsoft .NET Framework has largely superseded COM as the technology of choice for Windows application-level development, COM has had a lasting impact on the Windows development landscape, and it's far from a dead technology as it continues to be the foundation for many components (for example, the Windows shell UI uses COM extensively) and even some of the newest technologies in the Windows development landscape. Even on a Windows installation with no other additional applications, you will still find thousands of COM class identifiers (CLSID for short) describing existing COM classes (servers) in the COM registry catalog, as shown in Figure 1-7.

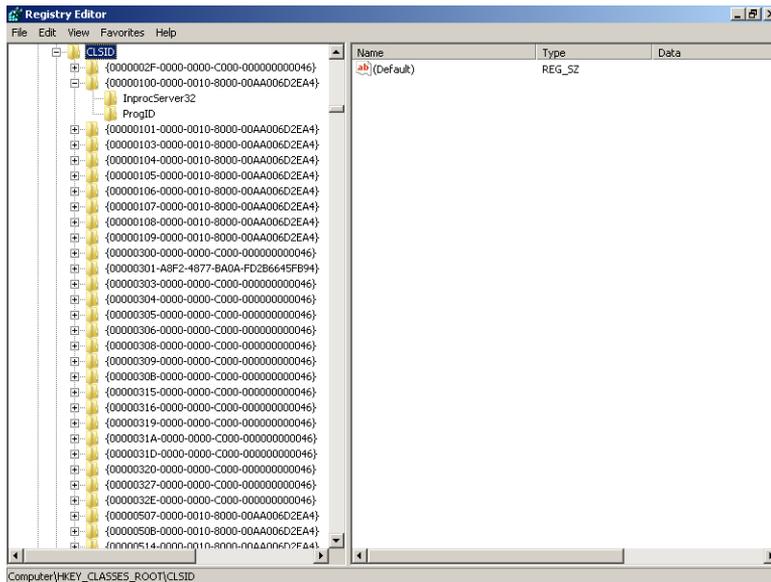


FIGURE 1-7 COM CLSID hive in the Windows registry.

More recently, the new model unveiled by Microsoft for developing touch-capable, Windows runtime (WinRT) applications in its upcoming version of Windows also uses COM for its core binary compatibility layer. In many ways, COM remains required knowledge if you really want to master the intricacies of Windows as a development platform. In addition, its design patterns and programming model have an educational value of their own. Even if you are never going to write a COM object (server), the programs and scripts you write often use existing COM objects, either explicitly or

indirectly via API calls. So, knowing how COM works is particularly useful in debugging situations that involve such calls.

COM developers usually interact primarily with language-level features and tools when writing or consuming COM servers. However, debugging COM failures also requires knowledge of how the system implements the COM specification. In that sense, there are a few different aspects to the COM landscape that can all come into play during COM debugging investigations:

- **The COM “Library”** This is essentially Microsoft’s system implementation of the COM binary specification. The COM library consists primarily of the COM run-time code and Win32 APIs (*CoInitialize*, *CoCreateInstance*, and others) that ship as part of the *ole32.dll* system module. This run-time code also uses support provided by two Windows services collectively referred to as the *COM service control manager (COM SCM)*. These services are the *RpcSs* service, which runs with *NetworkService* privileges, and the *DComLaunch* service, which runs with *LocalSystem* privileges.
- **COM language tools** These are the source-level compilers and tools to support COM’s binary specification. This includes the interface definition language (IDL) compiler (MIDL), which allows COM classes and interfaces to be consumed inside C/C++ programs, and the binary type library importer and exporter tools, which enable cross-language interoperability in COM. Note that these tools do not ship with the OS, but come as part of developer tools such as the Windows SDK.
- **COM frameworks** These are frameworks that make it easier for developers to write COM components that conform to the COM binary specification. An example is the Microsoft C++ Active Template Library (ATL).

Writing COM Servers

The COM specification places several requirements on developers writing COM objects and their hosting modules. Writing a COM object in C++ entails, at the very least, declaring its published interfaces in an IDL file and implementing the standard *IUnknown* COM interface that allows reference counting of the object and enables clients to negotiate contracts with the server. A class factory object—a COM-style support object that implements the *IClassFactory* COM interface but doesn’t need to be published to the registry—must also be written for each CLSID to create its object instances.

On top of all of this, the developer is also required to implement the hosting module (DLL or EXE) so that it also conforms to all the other requirements of the COM specification. For example, in the case of a DLL module, a C-style function (*DllGetClassObject*) that returns a pointer to the class factory of CLSIDs hosted by the module must be exported for use by the COM library. For an executable module, the COM library can’t simply call an exported function, so *ole32!CoRegisterClassObjects* must be called by the server executable itself when it starts up in order to publish its hosted COM class factories and make COM aware of them. Yet another requirement for DLL COM modules is to implement reference counting of their active objects and export a C-style function (*DllCanUnloadNow*) so that the COM library knows when it’s safe to unload the module in question.

Microsoft realized that this is a lot to ask of C++ COM developers and introduced the Active Template Library (ATL) to help simplify writing COM server modules and objects in the C++ language. Although the majority of C++ COM developers in Windows use ATL to implement their COM servers, keep in mind that you can also write COM objects and their hosting modules without it (if you feel inclined to do so). In fact, ATL ships with the source code of its template classes, so you can study those implementation headers and see how ATL implements its functionality on top of the COM services provided by the COM library in the operating system. As you'll see after looking at that source code, ATL takes care of much of the heavy lifting and boilerplate code for writing COM servers so that you don't have to. This allows you to focus on writing the code that implements your business logic without much of the burden imposed by the necessary COM model plumbing.

Managing Module Lifetimes in COM Servers

To give you a practical scenario so that you get a better feel for the type of features ATL provides, consider the problem of managing the lifetime of COM servers. As mentioned earlier, COM servers implement reference counting to determine when their hosting server module (DLL or EXE) can be unloaded safely. In the COM ATL library, for example, every time a new COM object is created within a COM module (*CAtlDllModule*), a global (per-module) integer variable counting the number of live objects implemented by that module is incremented by one. This special reference count is often referred to as the COM module's *lock count*. When the last reference to the object is released, the lock count is decremented by one. In addition, a client is also allowed to put a lock on the server (that is, increment its global lock count by one) by calling the *LockServer* method of the *IClassFactory* standard COM interface (implemented by the *CComClassFactory* class in ATL).

In the case of DLL COM servers, this essentially provides a simple implementation to ref-count the COM module so that when its *DllCanUnloadNow* function is called, the module is able to see the number of outstanding client references it has (the lock count) and, when that number drops to zero, report to the COM library that it's free to unload the DLL if it wants to. A host process is able to force this cleanup at any time by calling the *ole32!CoFreeUnusedLibraries* COM runtime function, which internally invokes the exported *DllCanUnloadNow* functions from all in-process DLL COM server modules loaded in the client process and unloads the modules that return TRUE.

The lock count is also used by COM EXE servers. In the case of ATL, for example, COM EXE modules (*CAtlExeModule*) use an extra thread that periodically checks their lock count and shuts the process down after a certain period of inactivity after its lock count drops to zero. You can see this logic in the ATL implementation headers, as in the *MonitorShutdown* callback shown in the following listing.

```
//  
// c:\ddk\7600.16385.1\inc\atlbase.h  
//  
void MonitorShutdown() throw()  
{
```

```

while (1)
{
    ::WaitForSingleObject(m_hEventShutdown, INFINITE);
    DWORD dwWait = 0;
    do
    {
        m_bActivity = false;
        dwWait = ::WaitForSingleObject(m_hEventShutdown, m_dwTimeOut);
    } while (dwWait == WAIT_OBJECT_0);

    if (!m_bActivity && m_nLockCnt == 0)
    {
        ::CoSuspendClassObjects();
        if (m_nLockCnt == 0)
            break;
    }
}
::CloseHandle(m_hEventShutdown);
::PostThreadMessage(m_dwMainThreadID, WM_QUIT, 0, 0);
}

```

Note that COM as an OS component and platform leaves implementing this module reference counting scheme up to the developer writing the COM server. Fortunately, however, the good folks at Microsoft who wrote the ATL framework took care of this on behalf of all the Windows developers writing COM servers using that framework.

Consuming COM Objects

Communication between COM clients and servers is a two-step process:

- **COM activation** This is the first step in the communication, where the COM library locates the class factory of the requested CLSID. This is done by first consulting the COM registry catalog to find the module that hosts the implementation of the COM server. The COM client code initiates this step with a call to either the *ole32!CoCreateInstance* or *ole32!CoGetClassObject* Win32 API call. Note that *ole32!CoCreateInstance* is simply a convenient wrapper that first calls *ole32!CoGetClassObject* to obtain the class factory, and then invokes the *IClassFactory::CreateInstance* method implemented by that class factory object to finally create a new instance of the target COM server.
- **Method invocations** After the COM activation step retrieves a proxy or direct pointer to the COM class object, the client can then query the interfaces exposed by the object and directly invoke the exposed methods.

A key point to understand when consuming COM objects in your code is the potential involvement of the COM SCM behind the scenes to instantiate the hosting module for the COM object and its corresponding class factory during the COM activation step. This is done because COM clients sometimes need to activate out-of-process servers in different contexts, such as when the COM server object needs to run in processes with higher privileges or in different user sessions, which requires the participation of a broker process that runs with higher privileges (the *DComLaunch* service). Another

reason is that the *RpcSs* Windows service also handles cross-machine COM activation requests (the DCOM case) and implements the communication channel in a way that's completely transparent to both COM clients and servers. It's especially important to understand this involvement during debugging investigations of COM activation failures. Once the COM activation sequence retrieves the requested class factory, however, the COM client is then able to directly invoke COM methods published by the server class without any involvement on the part of the COM SCM. Figure 1-8 summarizes these steps and the key components involved during the COM activation sequence.

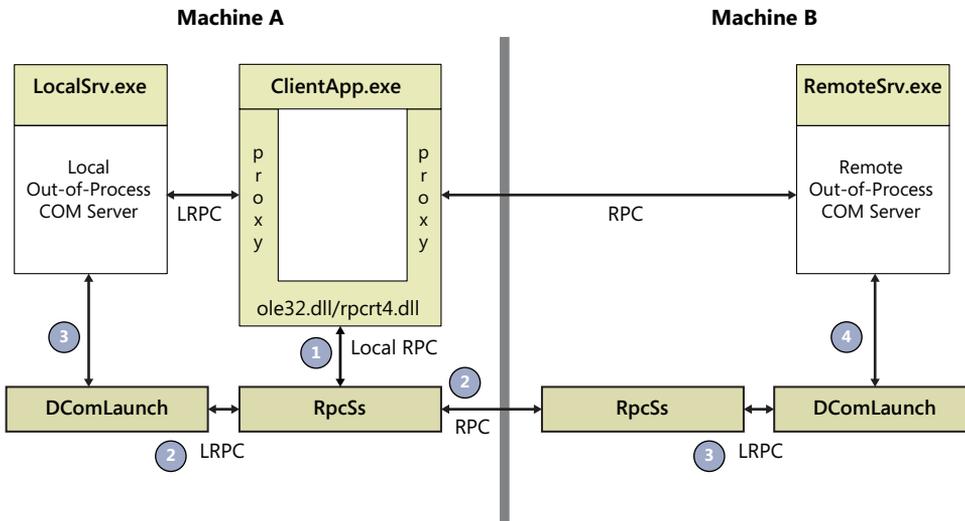


FIGURE 1-8 COM activation components.

The nice thing about the COM model, as mentioned earlier in this section, is that the client doesn't need to know where the COM server implementation lives, or even the language (C++, Microsoft Visual Basic, Delphi, or other) in which it was written (provided it has access to a type library or a direct virtual table layout that describes the COM types it wants to consume). The only thing that the client needs to know is the CLSID of the COM object (a GUID), after which it can query for the supported interfaces and invoke the desired methods. COM in the OS provides all the necessary "glue" for the client/server communication, provided the COM server was written to conform to the COM model. In particular, COM supports accessing the following COM server types using the same consistent programmatic model:

- **In-process COM servers** The hosting DLL module is loaded into the client process address space, and the object is invoked through a pointer returned by the COM activation sequence. This pointer can be either a direct virtual pointer or sometimes a proxy, depending on whether the COM runtime needs to be invoked to provide additional safeguards (such as thread safety) before invoking the methods of the COM server.
- **Local/remote out-of-process COM servers** For local out-of-process COM servers, local RPC is used as the underlying interprocess communication protocol, with ALPC as the actual low-level foundation. For remote COM servers (DCOM), the RPC communication protocol is used

for the intermachine communication. In both cases, the proxy memory pointer that is returned to the client application from the COM activation sequence takes care of everything that's required to accomplish COM's promise of transparent remoting. Figure 1-9 illustrates this aspect.

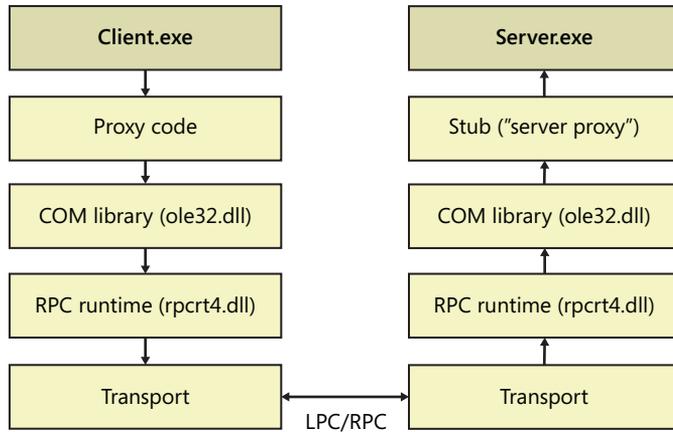


FIGURE 1-9 Out-of-process COM method invocations.

The CLR (.NET) Layer

Like COM, the .NET Framework is also a user-mode, object-oriented platform that enables developers to write their programs in their language of choice (C#, Microsoft Visual Basic .NET, C++/CLI, or other). However, .NET takes another leap and has those programs run under the control of an execution engine environment that provides additional useful features, such as type safety and automatic memory management using garbage collection. This execution engine environment is called the *Common Language Runtime (CLR)* and is in many ways the core of the .NET platform. Understanding this layer is often helpful when debugging applications built on top of the various .NET class libraries and technologies (ASP.NET, WCF, WinForms, WPF, Silverlight, and so on).

The CLR runtime is implemented as a set of native DLLs that get loaded into the address space of every .NET executable, though the core execution engine DLL decides when to load the other DLL dependencies. Because of their reliance on this execution environment, .NET modules (also called *assemblies*) are said to be *managed*, as opposed to the unmanaged native modules that execute in the regular user-mode environment. The same user-mode process can host both managed and unmanaged modules interoperating with each other, as will be explained shortly in this section.

Programs in .NET are not compiled directly into native assembly code, but rather into a platform-agnostic language called the *Microsoft .NET Intermediate Language* (usually referred to as MSIL, or simply IL). This IL is then lazily (methods are compiled on first use) turned into assembly instructions by a component of the execution engine called the *Just-in-Time .NET compiler*, or *JIT*.

.NET Side-by-Side Versioning

One of the issues that plagued software development in Windows prior to the introduction of the .NET Framework was the fact that new DLL versions sometimes introduced new behaviors, breaking existing software often through no fault of the application developer. There was no standard way to strongly bind applications to the version of the DLLs that they were tested against before they got released. This is known as the “DLL hell” issue. COM made the situation better by at least ensuring binary compatibility: instead of C-style exported DLL functions simply disappearing or altering their signatures from underneath their consumers (resulting in crashes!), COM servers were able to clearly version their interfaces, allowing COM clients to query the interfaces that they were tested against and providing the safety of this extra level of indirection.

The .NET Framework takes the idea of strong binding and versioning one step further by ensuring that .NET programs are always run against the version of the CLR that they were compiled to target or, alternatively, the version specified in the application’s configuration file. So, new versions of the .NET Framework are installed side by side with older ones instead of replacing them. The exception to this is a small shim DLL called *mscorlib.dll* that’s installed to the *system32* directory (on 64-bit Windows; a 32-bit version is also installed to the *SysWow64* directory) and that always matches the newest .NET Framework version present on the machine. This works because newer versions of *mscorlib.dll* are backward compatible with previous versions of the CLR. For example, if both .NET versions 4.0 and 2.0 are installed on the machine, the *mscorlib.dll* module in the *system32* directory will be the one that was installed with the CLR 4.0 release.

.NET Executable Programs Load Sequence

The IL assemblies produced by the various .NET compilers also follow the standard Windows PE (Portable Executable) format and are just special-case native images from the perspective of the OS loader, except they have a marker in their PE header to indicate they are managed code binaries that should be handled by the .NET CLR. When the Windows module loader code in *ntdll.dll* detects the existence of a CLR header in the executable PE image it is about to launch, control is immediately transferred to the native CLR entry point (*mscorlib!_CorExeMain*), which then takes care of finding and invoking the managed IL entry point in the image.



Note To support earlier OS versions (more specifically, Windows 98/ME and earlier service packs of Windows 2000 and Windows XP), which were not aware of the CLR header because the .NET Framework hadn’t shipped at the time those operating systems were released, managed PE images also had to have a regular native entry point consisting of a very thin stub (*jmp* instruction) that simply invokes the CLR’s main entry point method (*mscorlib!_CorExeMain*). Fortunately, the Windows releases supported by .NET are now natively capable of loading managed code modules, and this stub is no longer strictly required.

Note that the OS module loader doesn't really know which version of the CLR should be loaded for the managed image. This is the role of the *mscorlib.dll* native shim DLL, which determines the correct version of the CLR to load. For CLR v2 programs, for example, the execution engine DLL loaded by *mscorlib.dll* is *mscorwks.dll*, while for CLR v4 the execution engine implementation resides inside the *clr.dll* module. Once the CLR execution engine DLL for the target version is loaded, it pretty much takes control and becomes responsible for the runtime execution environment, ensuring type safety and automatic memory management (garbage collection), invoking the JIT compiler to convert IL into native assembly code on-demand, performing security checks, hosting the native implementation for several key threading and run-time services, and so on. The CLR is also responsible for loading the managed assemblies hosting the object types referenced by the application's IL, including the core library that implements some of the most basic managed types (*mscorlib.dll* .NET assembly). Figure 1-10 illustrates these steps.

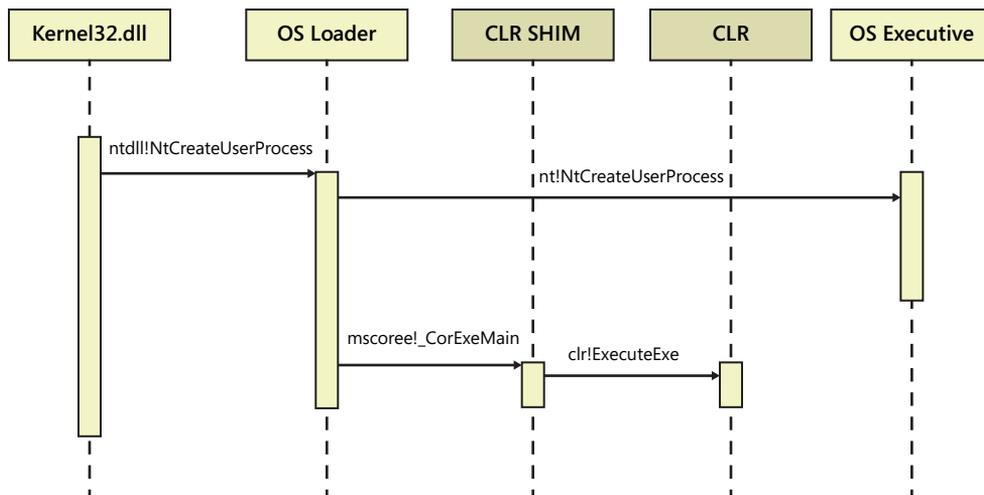


FIGURE 1-10 Loading steps for .NET executable programs.

Interoperability Between Managed and Native Code

When the first version of the .NET Framework was introduced back in early 2002, it had to be able to consume existing native code (both C-style Win32 APIs and COM objects) to ensure that the transition to managed code programming was seamless for Windows developers. Interoperability between managed and unmanaged code is possible fortunately, thanks in large part to two CLR mechanisms: *P/Invoke* and *COM Interop*. *P/Invoke* (the *DllImport* .NET method attribute) is used to invoke C-style unmanaged APIs such as those from the Win32 API set, and *COM Interop* (the *ComImport* .NET class attribute) can be used to invoke any existing classic COM object implemented in unmanaged code.

Managed/native code interoperability presents a few technical challenges, however. The biggest challenge is that the garbage collection scheme that the CLR uses for automatic memory management requires it to manage objects in the managed heap so that it is able to periodically clear

defunct (nonrooted) objects, and also so that it can reduce heap fragmentation when dead objects are collected. However, when transitioning to run native code as part of the same process address space, it's often necessary to share managed memory with the native call (in the form of function parameters, for instance). Because the CLR garbage collector is free to move these managed objects around when it decides to perform a garbage collection, and because the garbage collector is completely unaware of the operations that the native code might attempt, it could end up moving the shared managed objects around, causing the native code to access invalid memory. The CLR execution engine takes care of these technical intricacies by marshaling function parameters and pinning managed objects in memory during the managed to unmanaged transitions.

Conversely, you can also consume code that you develop using .NET from your native applications using the COM Interop facilities provided by the CLR. The C/C++ native code is able to consume the types published by a .NET assembly (by means of the *ComVisible* .NET attribute) using their type library, in the same fashion that native COM languages are able to consume types from different languages. The .NET Framework ships with a tool called *regasm.exe*, which can be used to easily generate type libraries for the COM types in a .NET assembly. The .NET Framework and Windows SDKs also include a development tool, called *tlbexp.exe*, that's able to do the same thing.

The CLR shim DLL (*mscorlib.dll*) again plays a key role in this reverse COM Interop scenario because it's the first native entry point to the CLR during the COM activation. This shim then loads the right CLR execution engine version, which then loads the managed COM types as they get invoked by the native code. This extends the functionality provided by the COM library in the OS without it having to know about the intricacies of managed code. During the COM activation sequence that the native application initiates, the COM library ends up invoking the standard *DllGetClassObject* method exported from *mscorlib.dll*. If you used *regasm.exe* to generate the type library for the C# COM types, *mscorlib.dll* also would've been added to the registry as the *InProcServer32* for all the managed COM classes hosted by the .NET DLL assembly. The CLR shim DLL then forwards the call to the CLR execution engine, which takes care of implementing the class factory and standard native COM interfaces on behalf of the managed COM types.

Microsoft Developer Tools

Microsoft typically releases supporting tools (compilers, libraries, and the like) for developers to write code that targets its technologies. These releases are referred to as *development kits*. For example, there is a software development kit (SDK) for developers of Windows Phone applications, a .NET SDK that contains tools to write and sign code for the .NET Framework, an Xbox development kit (XDK) for game developers, and so on. Many of these development kits, including the .NET and Windows Phone SDKs, are available as free downloads from the Microsoft Download Center at <http://www.microsoft.com/downloads>.

The Windows team at Microsoft also ships two important software development kits that include many of the tools presented in this book: more specifically, the Windows Driver Development Kit

(DDK), which contains the build environment used to compile all the native C++ code samples in the book's companion source code, and the Windows Software Development Kit, which contains the Windows debuggers and Windows Performance Toolkit. Both of these development kits are free and available for download from the Microsoft Download Center.

The Windows DDK (WDK)

Each release of Windows is accompanied by a driver development kit targeted for use by Windows driver developers, which contains the headers, libraries, and tools needed for building drivers, as well as several code samples for writing WDM, KMDF, and UMDF drivers.

One of the most useful features included with this kit is a full-blown build and development environment that can be used not only for driver development but for any kind of C/C++ development. It includes C/C++ compilers and many other native development frameworks, including the STL (Standard C++ Template Library) and ATL (Active Template Library for building COM servers) template libraries and their respective implementation headers.

The native C++ code samples from the companion source code use a small portion of ATL that provides support for smart pointers and basic string and collection operations (arrays, hash tables, and so on). Although ATL also comes with the Microsoft Visual Studio suite, the build environment of the DDK was chosen for this book's companion source code so that readers without Visual Studio can follow the case studies and experiments presented in this book.

The Windows SDK

Another important development kit shipped to support new Windows releases is the Windows SDK. Microsoft sometimes ships more than one SDK version per major Windows release: for example, versions 7.0 and 7.1 of the SDK target Windows 7 developers, with version 7.1 bringing many improvements to some of the key tools covered in this book.

The Windows SDK contains useful documentation and samples for building applications on Windows, as well as the public (official) header files and import libraries that are required for compiling your native Windows applications. In addition, the Windows SDK also contains two of the main debugging and tracing tools covered in this book—namely, the Windows debuggers package and Windows Performance Toolkit.

Step-by-step instructions for how to acquire and install those SDK tools will be provided when they're introduced in the following parts of this book so that they're closer to where you will end up needing them.

Summary

This chapter introduced some common terminology and also served as a very short and, admittedly, fast-paced introduction to important layers (Kernel, Executive, NTDLL, Win32, COM/.NET) in the Windows architecture and software development landscape. The following points are also worth remembering as you read the rest of this book:

- Server and client Windows releases share a common kernel and follow a relatively similar release schedule, so it's important to know the client variant of each server release and vice versa. When this book states that a certain kernel feature is added in Windows 7, for example, it's also implied that the Windows Server 2008 R2 kernel has the same capability.
- The CPU architecture of your OS (*x86* or *x64*) is always important to know when performing debugging and tracing experiments.
- When studying a new API set or platform feature, you should try to piece together an architectural diagram in your mind and understand where that new feature fits relative to the existing OS layers and development frameworks. This comes in handy when you later need to debug or trace the code in your software that uses the feature in question.
- When analyzing development frameworks and whether to use them in building your software, you should also try to understand how they are implemented on top of the built-in OS services and what additional functionality they provide. This should help you make an informed choice as to what frameworks fit your scenario best and whether taking the dependency is worth the productivity benefits you would gain.

Many of the concepts discussed in this chapter will be revisited in practical debugging and tracing situations. If you are like me, you will find that these topics start hitting a lot closer to home, so to speak, once you get into the habit of using debugging and tracing to confirm your understanding of the theoretical background. So, don't hesitate to come back and consult this chapter again as you run the debugging and tracing explorations proposed in this book. You might want to double-check some of the theory that has been covered here.

How Windows Debuggers Work

In this chapter

User-Mode Debugging	85
Kernel-Mode Debugging	98
Managed-Code Debugging	103
Script Debugging	112
Remote Debugging	116
Summary	123

This chapter explains how different types of debuggers work in Microsoft Windows. If you know these architectural foundations, many debugger concepts and behaviors suddenly start making sense. For example, this chapter explains why certain debugger commands and features work only in user-mode or kernel-mode debugging. You'll also dive into the architecture of managed-code debugging and discover why .NET source-level debugging isn't currently supported by the Windows debuggers.

Following that discussion, you'll learn how the architecture of script debugging relates to that of .NET debugging. With HTML5 and JavaScript becoming more prevalent development technologies for rich client user interfaces, script debugging is likely to garner even more attention from Windows developers in the future. This chapter concludes with a section that explains remote debugging and the key concepts that drive its architecture.

User-Mode Debugging

A user-mode debugger gives you the ability to inspect the memory of the target program you're trying to debug, as well as the ability to control its execution. In particular, you want to be able to set breakpoints and step through the target code, one instruction or one source line at a time. These basic requirements drive the design of the native user-mode debugging architecture in Windows.

Architecture Overview

To support controlling the target in user-mode debugging, the Windows operating system (OS) has an architecture based on the following principles:

- When important debug events, such as new module loads and exceptions, occur in the context of a process that's being debugged by a user-mode debugger, the OS generates message notifications on behalf of the target and sends them to the debugger process, giving the debugger program a chance to either handle or ignore those notifications. During each notification, the target process blocks and waits until the debugger is done responding to it before it resumes its execution.
- For this architecture to work, the native debugger process must also implement its end of the handshake, so to speak, and have a dedicated thread to receive and respond to the debug events generated by the target process.
- The interprocess communication between the two user-mode programs is based on a debug port kernel object (owned by the target process), where the target queues up its debug event notifications and waits on the debugger to process them.

This generic interprocess communication model is sufficient to handle all the requirements for controlling the target in a user-mode debugging session, providing the debugger with the capability to respond to code breakpoints or single-step events, as illustrated in Figure 3-1.

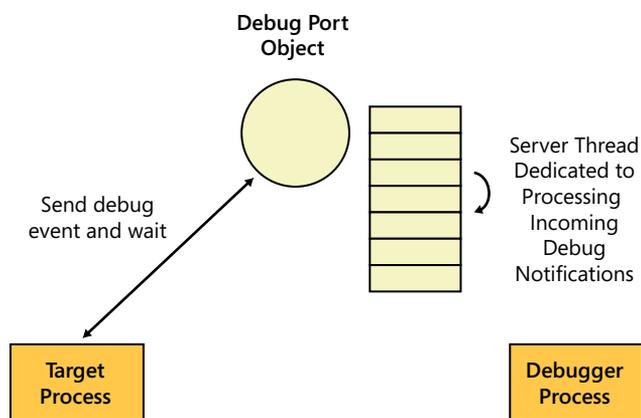


FIGURE 3-1 Native user-mode debugging architecture in Windows.

The other high-level requirement of user-mode debugging is for the debugger to be able to inspect and modify the virtual address space of the target process. This is necessary, for example, to be able to insert code breakpoints or walk the stacks and list the call frames in the threads of execution contained within the target process.

Windows provides facilities exposed at the Win32 API layer to satisfy these requirements, allowing any user-mode process to read and write to the memory of another process—as long as it has

sufficient privileges to do so. This system-brokered access is why you can debug only your own processes unless you're an administrator running in an elevated User Account Control (UAC) context with full administrative privileges (which include, in particular, the special *SeDebugPrivilege*). If you do have those privileges, you can debug processes from any other user on the system—including *LocalSystem* processes.

Win32 Debugging APIs

Debugger programs can implement their functionality and follow the conceptual model described in the previous section by using published APIs in the operating system. Table 3-1 summarizes the main Win32 functions used in Windows user-mode debuggers to achieve their requirements.

TABLE 3-1 Win32 API Support for User-Mode Windows Debuggers

Requirement	Win32 API Function	WinDbg Command(s)
Start a target process directly under the control of a user-mode debugger.	<i>CreateProcess</i> , with <i>dwCreationFlags</i> : <ul style="list-style-type: none"> ■ <code>DEBUG_PROCESS</code> ■ <code>DEBUG_ONLY_THIS_PROCESS</code> 	Ctrl+E UI shortcut or windbg.exe target.exe
Dynamically attach a user-mode debugger to an existing process.	<i>OpenProcess</i> , with at least the following <i>dwDesiredAccess</i> flags: <ul style="list-style-type: none"> ■ <code>PROCESS_VM_READ</code> ■ <code>PROCESS_VM_WRITE</code> ■ <code>PROCESS_VM_OPERATION</code> <i>DebugActiveProcess</i> , with the handle obtained in the previous step	F6 UI shortcut or windbg.exe -pn target.exe or windbg.exe -p [PID]
Stop debugging the target process, but without terminating it.	<i>DebugActiveProcessStop</i>	qd ("quit and detach")
Break into the debugger to inspect the target.	<i>DebugBreakProcess</i>	Ctrl+Break UI shortcut or Debug\Break menu action
Wait for new debug events.	<i>WaitForDebugEvent</i>	N/A
Continue the target's execution after a received debug event is processed.	<i>ContinueDebugEvent</i>	N/A
Inspect and edit the virtual address space of the target process.	<i>ReadProcessMemory</i> <i>WriteProcessMemory</i>	Dump memory (<i>dd</i> , <i>db</i> , and so on) Edit memory (<i>ed</i> , <i>eb</i> , and so on) Insert code breakpoints (<i>bp</i>) Dump a thread's stack trace (<i>k</i> , <i>kP</i> , <i>kn</i> , and so on)

With these Win32 APIs, a user-mode debugger can write the code in the thread that it uses to process the debug events it receives from the target using a loop like the one shown in the following listing (pseudo-code).

```
//
// Main User-Mode Debugger Loop
//
CreateProcess("target.exe", ..., DEBUG_PROCESS, ...);
while (1)
```

```

{
    WaitForDebugEvent(&event, ...);
    switch (event)
    {
        case ModuleLoad:
            Handle/Ignore;
            break;
        case TerminateProcess:
            Handle/Ignore;
            break;
        case Exception (code breakpoint, single step, etc...):
            Handle/Ignore;
            break;
    }
    ContinueDebugEvent(...);
}

```

When the debugger loop calls the *WaitForDebugEvent* Win32 API to check whether a new debug event arrived, the call internally transitions to kernel mode to fetch the event from the queue of the debug port object of the target process (the *DebugPort* field of the *nt!_EPROCESS* executive object). If the queue is found to be empty, the call blocks and waits for a new debug event to be posted to the port object. After the event is processed by the debugger, the *ContinueDebugEvent* Win32 API is called to let the target process continue its execution.

Debug Events and Exceptions

The OS generates several types of debug events when a process is being debugged. For instance, an event is generated for every module load, allowing the user-mode debugger to know when a new DLL is mapped into the address space of the target process. Similarly, an event is also raised when a new child process is created by the target process, enabling the user-mode debugging session to also handle the debug events from child processes if it wants to do so.

Debug events are similarly generated when any exceptions are raised in the context of the target process. As you'll shortly see, code breakpoints and single-stepping are both internally implemented by forcing an exception to be raised in the context of the target application, which means that those events can also be handled by the user-mode debugger just like any other debug events. To better understand this type of debug event, a quick overview of exception handling in Windows is in order.

.NET, C++, and SEH Exceptions

Two categories of exceptions exist in Windows: language-level or framework-level exceptions, such as the C++ or .NET exceptions, and OS-level exceptions, also known as *Structured Exception Handling* (SEH) exceptions. Both Microsoft Visual C++ and the .NET Common Language Runtime (CLR) use SEH exceptions internally to implement support for their specific application-level, exception-handling mechanisms.

SEH exceptions, in turn, can be separated into two categories: hardware exceptions are raised in response to a processor interrupt (invalid memory access, integer divide-by-zero, and so on), and software exceptions are triggered by an explicit call to the *RaiseException* Win32 API. Hardware exceptions are particularly important to the functionality of user-mode debuggers in Windows because they're also used to implement breakpoints and in single-stepping the target, two fundamental features of any debugger.

At the Visual C/C++ language level, the *throw* keyword used to throw C++ exceptions is translated by the compiler into a call implemented in the C runtime library, which ultimately invokes the *RaiseException* API. In addition, three keywords (*__try*, *__except*, and *__finally*) are also defined to allow you to take advantage of SEH exceptions and structure (hence the SEH name) your code so that you can establish code blocks to handle or ignore the SEH exceptions that get raised from within that code block. Despite this Visual C++ language support, it's important to realize that SEH is a Windows operating system concept and that you can use it with any language, as long as the compiler has support for it.

Unlike C++ exceptions, which can be raised with any type, SEH exceptions deal with only one type: *unsigned int*. Each SEH exception, whether triggered in hardware or software, is identified in Windows using an integer identifier—the *exception code*—that indicates the type of fault that triggered the exception (divide-by-zero, access violation, and so on). You can find many of the exception codes defined by the OS in the *winnt.h* Software Development Kit (SDK) header file. In addition, applications are also free to define their own custom exception codes, which is precisely what the C++ and .NET runtimes do for their exceptions.

Table 3-2 lists a few common exception codes that you'll see frequently during your debugging investigations.

TABLE 3-2 Common Windows SEH Exceptions and Their Status Codes

Exception Code	Description
<i>STATUS_ACCESS_VIOLATION (0xC0000005)</i>	Invalid memory access
<i>STATUS_INTEGER_DIVIDE_BY_ZERO (0xC0000094)</i>	Arithmetic divide-by-zero operation
<i>STATUS_INTEGER_OVERFLOW (0xC0000095)</i>	Arithmetic integer overflow
<i>STATUS_STACK_OVERFLOW (0xC00000FD)</i>	Stack overflow (running out of stack space)
<i>STATUS_BREAKPOINT (0x80000003)</i>	Raised in response to the debug break CPU interrupt (interrupt #3 on x86 and x64)
<i>STATUS_SINGLE_STEP (0x80000004)</i>	Raised in response to the single-step CPU interrupt (interrupt #1 on x86 and x64)

SEH Exception Handling in the User-Mode Debugger

When an exception occurs in a process that's being debugged, the user-mode debugger gets notified by the OS exception dispatching code in *ntdll.dll* before any user exception handlers defined in the target process are given a chance to respond to the exception. If the debugger chooses not to handle

this *first-chance* exception notification, the exception dispatching sequence proceeds further and the target thread is then given a chance to handle the exception if it wants to do so. If the SEH exception is not handled by the thread in the target process, the debugger is then sent another debug event, called a *second-chance* notification, to inform it that an unhandled exception occurred in the target process.

Figure 3-2 summarizes this OS exception dispatching sequence, specifically when a user-mode debugger is connected to the target process.

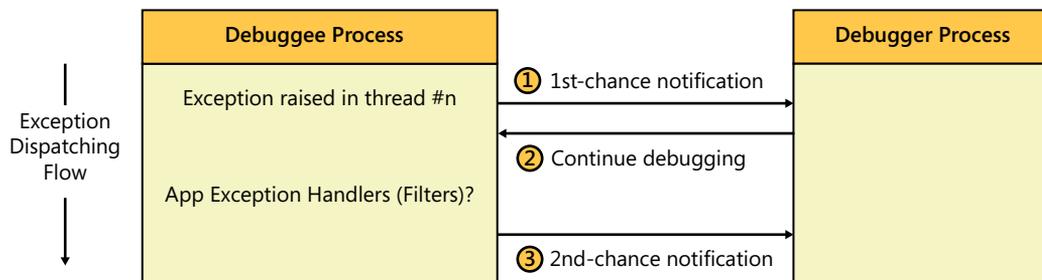


FIGURE 3-2 SEH exceptions and debug event notifications.

First-chance notifications are a good place for the user-mode debugger to handle exceptions that should be invisible to the code in the target process, including code breakpoints, single-step debug events, and the break-in signal. The sections that follow describe these important mechanisms in more detail.

Unlike first-chance notifications, which for user exceptions are simply logged to the debugger command window by default, the user-mode debugger always stops the target in response to a second-chance exception notification. Unhandled exceptions are always reason for concern because they lead to the demise of the target process when no debuggers are attached, which is why the user-mode debugger breaks in when they occur so that you can investigate them. You can see this sequence in action using the following program from the companion source code, which simply throws a C++ exception with a string type. For more details on how to compile the companion source code, refer to the procedure described in the Introduction of this book.

```
//  
// C:\book\code\chapter_03\BasicException>main.cpp  
//  
int  
__cdecl  
wmain()  
{  
    throw "This program raised an error";  
    return 0;  
}
```

When you run this program under the WinDbg user-mode debugger, you see the debugger receive two notifications from the target: the first-chance notification is logged to the debugger command window, while the second-chance notification causes the debugger to break in, as illustrated in the following debugger listing. Notice also how the *throw* keyword used to raise C++ exceptions ends up getting translated into a call to the C runtime library (the *msvcrt!_CxxThrowException* function call in the following listing), which ultimately invokes the *RaiseException* Win32 API to raise an SEH exception with the custom C++ exception code.

```

0:000> vercommand
command line: '"c:\Program Files\Debugging Tools for Windows (x86)\windbg.exe"
c:\book\code\chapter_03\BasicException\objfre_win7_x86\i386\BasicException.exe'
0:000> .symfix
0:000> .reload
0:000> g
(aa8.1fc0): C++ EH exception - code e06d7363 (first chance)
(aa8.1fc0): C++ EH exception - code e06d7363 (!!! second chance !!!)
...
KERNELBASE!RaiseException+0x58:
75dad36f c9          leave
0:000> k
ChildEBP RetAddr
000ffb60 75fd359c KERNELBASE!RaiseException+0x58
000ffb98 00cb1204 msvcrt!_CxxThrowException+0x48
000ffb9c 00cb136d BasicException!wmain+0x1b
[c:\book\code\chapter_03\basicexception\main.cpp @ 7]
000ffbf0 76f9ed6c BasicException!__wmainCRTStartup+0x102
000ffbf4 779c377b kernel32!BaseThreadInitThunk+0xe
000ffc3c 779c374e ntdll!_RtlUserThreadStart+0x70
000ffc54 00000000 ntdll!_RtlUserThreadStart+0x1b
0:000> $ Quit the debugging session
0:000> q

```

The Break-in Sequence

User-mode debuggers can intervene at any point in time and freeze the execution of their target process so that it can be inspected by the user—an operation referred to as *breaking into* the debugger. This is achieved by using the *DebugBreakProcess* API, which internally injects a remote thread into the address space of the target process. This “break-in” thread executes a debug break CPU interrupt instruction (*int 3*). In response to this interrupt, an SEH exception is raised by the OS in the context of the break-in thread. As shown in the previous section, this sends the user-mode debugger process a first-chance notification, allowing it to handle this special debug break exception (code *0x80000003*, or *STATUS_BREAKPOINT*) and finally break in by suspending all the threads in the target process.

This is why the current thread context in the user-mode debugger after a break-in operation will be in this special thread, which isn’t a thread you’ll recognize as “yours” if you’re debugging your own target process. To see this break-in thread in action, start a new instance of *notepad.exe* under the WinDbg user-mode debugger, as shown in the following listing. If you’re running this experiment

on 64-bit Windows, you can execute it again exactly as shown here by using the 32-bit version of *notepad.exe* located under the *%windir%\SysWow64* directory on *x64* Windows.

```
0:000> vercommand
command line: '"c:\Program Files\Debugging Tools for Windows (x86)\windbg.exe" notepad.exe'
0:000> .symfix
0:000> .reload
Reloading current modules.....
0:000> ~
    0 Id: 1d90.1678 Suspend: 1 Teb: 7ffde000 Unfrozen
0:000> g
```

Using the *Debug\Break* menu action, break back into the debugger. You'll see that the current thread context is no longer thread #0 (the main UI thread in *notepad.exe*) but rather a new thread. As you can infer from the function name (*ntdll!DbgUiRemoteBreakin*) on the call stack that you obtain by using the *k* command, this is the remote thread that was injected by the debugger into the target address space in response to the break-in request.

```
(1938.1fb0): Break instruction exception - code 80000003 (first chance)
...
ntdll!DbgBreakPoint:
7799410c cc          int     3
0:001> ~
    0 Id: 1d90.1678 Suspend: 1 Teb: 7ffde000 Unfrozen
    1 Id: 1d90.17f0 Suspend: 1 Teb: 7ffdd000 Unfrozen
0:001> k
ChildEBP RetAddr
00a4fecc 779ef161 ntdll!DbgBreakPoint
00a4fefc 75e9ed6c ntdll!DbgUiRemoteBreakin+0x3c
00a4ff08 779b37f5 kernel32!BaseThreadInitThunk+0xe
00a4ff48 779b37c8 ntdll!_RtlUserThreadStart+0x70
00a4ff60 00000000 ntdll!_RtlUserThreadStart+0x1b
```

In addition, using the *uf* debugger command to disassemble the current function shows that this thread was executing an *int 3* CPU interrupt instruction just before the debugger got sent the first-chance notification for the debug break exception.

```
0:001> uf .
ntdll!DbgBreakPoint:
7799410c cc          int     3
7799410d c3          ret
```

To see the actual threads in the target process, you can use the *~*k* command to list the call stacks for every thread in the target. You can also use the *s* command to change ("switch") the current thread context in the debugger to one of those threads, as illustrated in the following listing.

```
0:001> $ Switch over to thread #0 in the target
0:001> ~0s
```

```

0:000> k
ChildEBP RetAddr
0019f8e8 760fcde0 ntdll!KiFastSystemCallRet
0019f8ec 760fce13 USER32!NtUserGetMessage+0xc
0019f908 0085148a USER32!GetMessageW+0x33
0019f948 008516ec notepad!WinMain+0xe6
0019f9d8 76f9ed6c notepad!_initterm_e+0x1a1
0019f9e4 779c377b kernel32!BaseThreadInitThunk+0xe
0019fa24 779c374e ntdll!__RtlUserThreadStart+0x70
0019fa3c 00000000 ntdll!_RtlUserThreadStart+0x1b
0:001> $ Terminate this debugging session...
0:001> q

```

Setting Code Breakpoints

Code breakpoints are also implemented using the *int 3* instruction. Unlike the break-in case, where the debug break instruction is executed in the context of the remote break-in thread, code breakpoints are implemented by directly overwriting the target memory location where the code breakpoint was requested by the user.

The debugger program keeps track of the initial instructions for each code breakpoint so that it can substitute them in place of the debug break instruction when the breakpoints are hit, and before the user is able to inspect the target inside the debugger. This way, the fact that *int 3* instructions are inserted into the target process to implement code breakpoints is completely hidden from the user debugging the program, as it should be.

This scheme sounds straightforward, but there is a catch: how is the debugger able to insert the *int 3* instruction before the execution of the target process is resumed (using the *g* command) after a breakpoint hit? Surely, the debugger can't simply insert the debug break instruction before the target's execution is resumed because the next instruction to execute is supposed to be the original one from the target and not the *int 3* instruction. The way the debugger solves this dilemma is the same way it is able to support single-stepping, which is by using the TF ("trap flag") bit of the EFLAGS register on *x86* and *x64* processors to force the target thread to execute one instruction at a time. This single-step flag causes the CPU to issue an interrupt (*int 1*) after every instruction it executes. This allows the thread of the breakpoint to execute the original target instruction before the debugger is immediately given a chance to handle the new single-step SEH exception—which it does by restoring the debug break instruction again, as well as by resetting the TF flag so that the CPU single-step mode is disabled again.

Observing Code Breakpoint Insertion in WinDbg

To conclude this section, you'll try a fun experiment in which you'll debug the user-mode WinDbg debugger! Armed with the background information from this section and the familiarity with using WinDbg commands that you've gained so far, you have all the tools to confirm what WinDbg does when a new code breakpoint is added by the user without having to take my word for it.

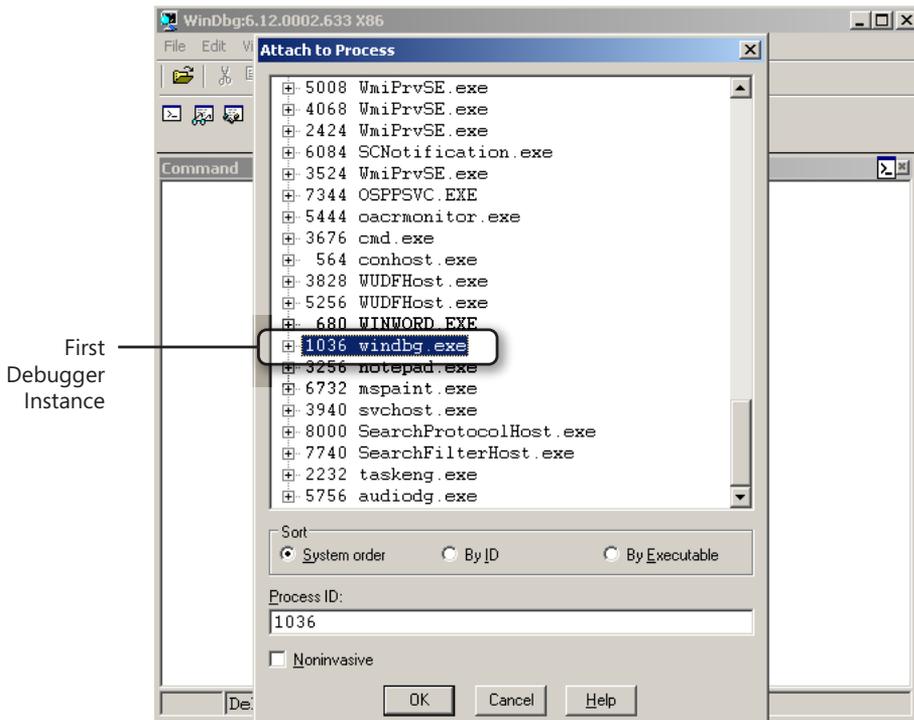


FIGURE 3-4 Debugging the debugger: the second WinDbg debugger instance.

In this new WinDbg instance, set a breakpoint at the `kernel32!WriteProcessMemory` API. As mentioned earlier in this chapter, this is the Win32 API used by user-mode debuggers to edit the virtual memory of their target processes.

```
0:002> $ Second WinDbg Session
0:002> .symfix
0:002> .reload
0:002> x kernel32!*writeprocessmemory*
75e51928 kernel32!_imp__WriteProcessMemory = <no type information>
75e520e3 kernel32!WriteProcessMemory = <no type information>
75eb959f kernel32!WriteProcessMemoryStub = <no type information>
0:002> bp kernel32!WriteProcessMemory
0:002> g
```

Now that you have this breakpoint in place, go back to the first `windbg.exe` instance and run the `g` command to let `notepad.exe` continue its execution.

```
0:000> $ First WinDbg Session
0:000> g
```

Notice that you immediately get a breakpoint hit inside the second `windbg.exe` instance, which is consistent with what you already learned in this chapter, because the first debugger tries to insert an `int 3` instruction into the `notepad.exe` process address space (corresponding to the `USER32!GetMessageW` breakpoint you added earlier).

```

Breakpoint 0 hit
kernel32!WriteProcessMemory:
...
0:001> $ Second WinDbg Session
0:001> k
ChildEBP RetAddr
0092edf4 58b84448 kernel32!WriteProcessMemory
0092ee2c 58adb384 dbgeng!BaseX86MachineInfo::InsertBreakpointInstruction+0x128
0092ee7c 58ad38ee dbgeng!LiveUserTargetInfo::InsertCodeBreakpoint+0x64
0092eeb8 58ad62f7 dbgeng!CodeBreakpoint::Insert+0xae
0092f764 58b67719 dbgeng!InsertBreakpoints+0x8c7
0092f7e8 58b66678 dbgeng!PrepareForExecution+0x5d9
0092f7fc 58afa539 dbgeng!PrepareForWait+0x28
0092f840 58afaa60 dbgeng!RawWaitForEvent+0x19
0092f858 00ebb6cf dbgeng!DebugClient::WaitForEvent+0xb0
0092f874 75e9ed6c windbg!EngineLoop+0x13f
0092f880 779b37f5 kernel32!BaseThreadInitThunk+0xe
0092f8c0 779b37c8 ntdll!__RtlUserThreadStart+0x70
0092f8d8 00000000 ntdll!_RtlUserThreadStart+0x1b

```

You can also check the arguments to the *WriteProcessMemory* API in the previous call stack by applying the technique described in Chapter 2, "Getting Started," where the stack pointer and saved frame pointer values were used to get the arguments to each call from the stack of the current thread. Remember that in the *__stdcall* calling convention, the stack pointer register value points to the return address at the time of the breakpoint, followed by the arguments to the function call. This means that the second DWORD value in the following listing represents the first parameter to the Win32 API call. The values you'll see will be different, but you can apply the same steps described here to derive the function arguments to this API call:

```

0:001> $ Second WinDbg Session
0:001> dd esp
0092edd4 58ce14a2 00000120 7630cde8 58d1b5d8
0092ede4 00000001 0092edf0 00000000 0092ee2c
...

```

In the documentation for the *WriteProcessMemory* Win32 API on the MSDN website at <http://msdn.microsoft.com/>, you'll see that it takes five parameters.

```

BOOL
WINAPI
WriteProcessMemory(
    __in HANDLE hProcess,
    __in LPVOID lpBaseAddress,
    __in_bcount(nSize) LPCVOID lpBuffer,
    __in SIZE_T nSize,
    __out_opt SIZE_T * lpNumberOfBytesWritten
);

```

The first of these parameters is the user-mode handle for the target process object (*hProcess*) that the debugger is trying to write to. You can use the value you obtained from the *dd* command with the *!handle* debugger extension command to confirm that it was indeed the *notepad.exe* process.

The *!handle* command also gives you the process ID (PID) referenced by the handle, which you can confirm is the PID of *notepad.exe* in the Windows task manager UI or, alternatively, by using the convenient *.tlist* debugger command, as demonstrated in the following listing.

```
0:001> $ Second WinDbg Session
0:001> !handle 120 f
Handle 120
Type                Process
GrantedAccess       0x12167b:
                    ReadControl,Synch
                    Terminate,CreateThread,VMOp,VMRead,VMWrite,DupHandle,SetInfo,QueryInfo
Object Specific Information
  Process Id 5964
  Parent Process 3736
0:001> .tlist notepad*
0n5964 notepad.exe
```

Next is the address that the debugger is trying to overwrite (*lpBaseAddress*). Using the *u* debugger command to disassemble the code located at that address, you can see that this second argument does indeed point to the *USER32!GetMessageW* API, which was the target location of the requested code breakpoint.

```
0:001> $ Second WinDbg Session
0:001> u 0x7630cde8
USER32!GetMessageW:
7630cde8 8bff          mov     edi,edi
7630cdea 55           push   ebp
7630cdeb 8bec        mov     ebp,esp
7630cded 8b5510      mov     edx,dword ptr [ebp+10h]
...
```

Finally, the third parameter (*lpBuffer*) is a pointer to the buffer that the debugger is trying to insert into this memory location. This is a single-byte buffer (as indicated by the value of the fourth argument, *nSize*, from the previous listing), representing the *int 3* instruction. On both *x86* and *x64*, this instruction is encoded using the single *0xCC* byte, as you can see by using either the *u* (“un-assemble”) or *db* (“dump memory as a sequence of bytes”) commands:

```
0:001> $ Second WinDbg Session
0:001> u 58d1b5d8
dbgeng!g_X86Int3:
58d1b5d8 cc          int     3
0:001> db 58d1b5d8
58d1b5d8 cc 00 00 00 00 00 00 00-00 00 00 00 00 00 00 00 .....
...
```

If you continue this experiment with the same breakpoints and type *g* again inside the second *windbg.exe* instance, you can similarly analyze the next hit of the *WriteProcessMemory* breakpoint and confirm that the initial byte from the *USER32!GetMessageW* function (*0x8b* in this case) is

surreptitiously restored as the *USER32!GetMessageW* breakpoint gets hit, right before the user is able to issue commands in the first debugger UI, as shown in the following listing.

```
0:001> $ Second WinDbg Session
0:001> g
Breakpoint 0 hit
0:001> k
ChildEBP RetAddr
0092f128 58b84542 kernel32!WriteProcessMemory
0092f160 58adb6a9 dbgeng!BaseX86MachineInfo::RemoveBreakpointInstruction+0xa2
0092f1a4 58ad3bcd dbgeng!LiveUserTargetInfo::RemoveCodeBreakpoint+0x59
0092f1ec 58ad6c0a dbgeng!CodeBreakpoint::Remove+0x11d
...
0:001> dd esp
0092f108 58ce14a2 00000120 7630cde8 00680cc4
0092f118 00000001 0092f124 00479ed8 0000000b
...
0:001> db 00680cc4
00680cc4 8b 00 00 00 00 00 00 00-00 00 00 00 00 00 00 00 .....
...
0:001> $ Terminate both debugging sessions now
0:001> q
```

Kernel-Mode Debugging

The high-level requirements for kernel-mode debugging are similar to those of user-mode debugging, including the ability to control the target (break in, single-step, set breakpoints, and so on) and also manipulate its memory address space. The difference in the case of kernel-mode debugging is that the target is the entire system being debugged.

Architecture Overview

Just like in the case of user-mode debugging, the Windows operating system also designed an architecture that answers the system-level needs of kernel debuggers. In the case of user-mode debugging, that support framework is built right into the OS kernel, where the debug port executive object provides the key to the interprocess communication channel between the debugger and target processes. In the case of kernel debugging, the kernel itself is being debugged, so support for the communication channel is built lower in the architectural stack. This is done using Hardware Abstraction Layer (HAL) extensions that implement the low-level transport layer of the communication channel between the host and target machines during kernel debugging.

There are different transport mediums you can use to perform kernel-mode debugging, and each one of them is implemented in its own transport DLL extension. In Windows 7, for example, *kdcorn.dll* is used for serial cables, *kd1394.dll* is used for FireWire cables, and *kdusb.dll* is used for USB 2.0 debug cables. These module extensions are loaded by the HAL very early during the boot process, when the target is enabled to support kernel-mode debugging. Because these modules sit very low in the architecture stack, they can't depend on higher-level OS kernel components that might not yet be fully loaded or otherwise turn out to be themselves in the process of being debugged. For that

reason, the KD transport extensions are fairly lightweight and interact directly with the hardware at the lowest possible level without taking any extra device driver dependencies, as demonstrated in Figure 3-5.

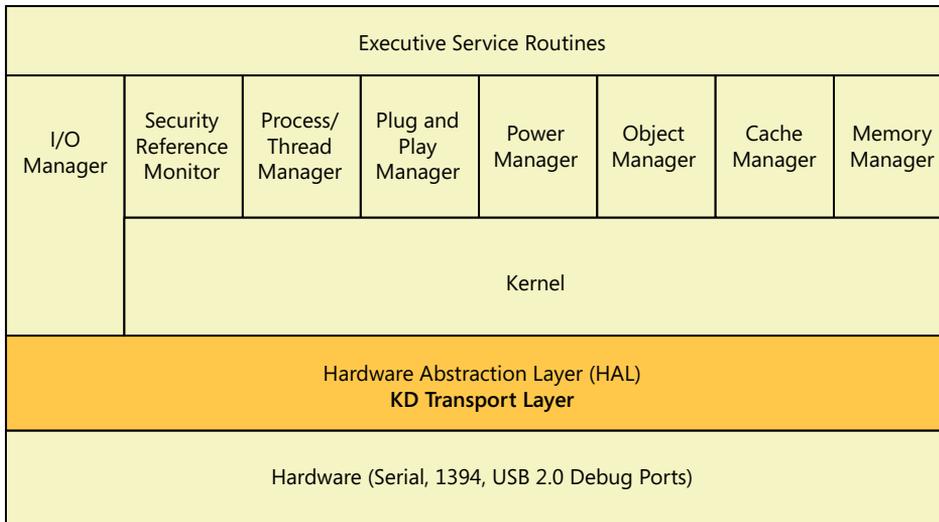


FIGURE 3-5 KD transport layer in the target OS.

If you disregard for a second how the debugger commands are transmitted from the kernel debugger to the target, the conceptual model for how the kernel on the target processes the commands sent by the kernel-mode debugger is quite similar to how debug events are processed by the user-mode debugger loop:

- The OS kernel periodically asks the transport layer (as part of the clock interrupt service routine) to check for break-in packets from the host debugger. When a new one is found, the kernel enters a break-in loop where it waits for additional commands to be received from the host kernel debugger.
- While the system on the target machine is halted, the break-in loop checks for any new commands sent by the host kernel debugger. This enables the kernel debugger to read register values, inspect or change memory on the target, and perform many other inspection and control commands while the target is still frozen. These send/receive handshakes are repeated until the host kernel debugger decides to leave the break-in state and the target is instructed to exit the debugger break-in mode and continue its normal execution again.
- In addition to explicit break-in requests, the kernel can also enter the break-in loop in response to exceptions that get raised by the target machine, which allows the debugger to intervene and respond to them. This generic handling of exceptions is again used to implement single-stepping and setting code breakpoints inside the target OS during kernel-mode debugging.

Setting Code Breakpoints

Knowing how code breakpoints are implemented during kernel-mode debugging is important so that you can understand situations when you fail to hit breakpoints you insert using the host kernel debugger. There are many similarities between how code breakpoints are internally implemented in user-mode and kernel-mode debugging, but there are also several important differences.

Like in the user-mode debugging case, code breakpoints are also inserted by overwriting the target virtual memory address with the debug break CPU instruction (*int 3*). When the target machine hits the inserted breakpoint, a CPU interrupt is raised and its OS interrupt handler is invoked. Where things diverge between user-mode and kernel-mode debugging is in how the handler dispatches the exception event to the host debugger. In the kernel-mode debugging case, the target OS is halted and enters the break-in send/receive loop, allowing the host debugger to handle the breakpoint by putting the initial byte back in the breakpoint's code location before entering the break-in state.

Another way that kernel debugging code breakpoints are different from their user-mode debugging counterparts is that they might refer to memory that has been paged out to disk on the target machine. In that case, the target simply handles the breakpoint command from the host debugger by registering the code breakpoint as being "owed." When the code page is later loaded into memory, the page fault handler (*nt!MmAccessFault*) in the kernel memory manager intervenes and inserts the breakpoint instruction to the global code page at that time, just as it would have done if the breakpoint had been in a memory location that wasn't paged out at the time of the debugger break-in.

Finally, because the same user-mode virtual memory address can point to different private code depending on the user-mode process context, code breakpoints inserted during kernel debugging are always interpreted relative to the current process context. This is a point that sometimes escapes developers who are new to kernel debugging because it isn't a concern in user-mode debugging. However, this is precisely the reason why you should always invasively switch the process context in the host kernel debugger to the target process before setting breakpoints in user-mode code relative to that process.

Single-Stepping the Target

Single-stepping the target in the host debugger is implemented using the same single-step CPU support and interrupt (*int 1*) that enables you to single-step the target process in a user-mode debugging environment. However, the fact that kernel-mode debuggers have global scope again introduces some interesting side effects you should be aware of so that you are better prepared to deal with them during your kernel-debugging experiments.

The most practical difference you'll see when you try single-stepping the target in a host kernel debugger is that execution sometimes seems to jump to other random code on the system and away from your current thread context. This happens when the thread quantum expires while stepping over a function call and the OS decides to schedule another thread on the processor. When that happens, it seems as if the code you're debugging just jumped to a random location. In reality, what happened is that the old thread got switched out and a new one is now running on the processor.

This usually happens whenever you step over a long function or a Win32 API call that causes the thread to enter a wait state (such as a *Sleep* call). Fortunately, when single-stepping in a host kernel debugger, the target OS not only enables the CPU trace flag but cleverly also finds the next call and inserts an additional debug break instruction at that memory location every time you single-step. This means that by letting the target machine “go” again (using the *g* command) after it seemed you had jumped to an unrelated code location, you break right back at the next call from the original thread (once its wait is satisfied and the thread gets scheduled to run again), which allows you to continue single-stepping the thread you were examining prior to the context switch.

Switching the Current Process Context

There are two ways to resolve symbols for the user-mode stacks of a process on the target machine of a kernel-debugging session. The first way, which you already used in Chapter 2, is to simply switch the current process view in the host debugger and reload the user-mode symbols for that process. This method’s main advantage is that it also works in live kernel-mode debugging, where it proves useful when you need to observe multiple user-mode processes during a debugger break-in. In the following live kernel-debugging session, the *.process* command is used with the */r* (“reload user-mode symbols”) and */p* (“target process”) options to illustrate this important approach. Make sure you start a new *notepad.exe* instance and that you use the values in bold text when you execute these commands because your values are likely to be different from the ones shown in this listing.

```
Tkd> !process 0 0 notepad.exe
PROCESS 874fa030 SessionId: 1 Cid: 14ac Peb: 7ffdf000 ParentCid: 1348
Image: notepad.exe
Tkd> .process /r /p 874fa030
Implicit process is now 874fa030
Loading User Symbols.....
Tkd> !process 874fa030 7
PROCESS 874fa030 SessionId: 1 Cid: 14ac Peb: 7ffdf000 ParentCid: 1348
Image: notepad.exe
    THREAD 86f6fd48 Cid 14ac.2020 Teb: 7ffde000 Win32Thread: ffb91dd8 WAIT ...
        85685be8 SynchronizationEvent
        ChildEBP RetAddr Args to Child
        9a99fb10 8287e65d 86f6fd48 807c9308 807c6120 nt!KiSwapContext+0x26
        9a99fb48 8287d4b7 86f6fe08 86f6fd48 85685be8 nt!KiSwapThread+0x266
        9a99fb70 828770cf 86f6fd48 86f6fe08 00000000 nt!KiCommitThreadWait+0x1df
        9a99fbe8 9534959a 85685be8 0000000d 00000001 nt!KewaitForSingleObject+0x393
        9a99fc44 953493a7 000025ff 00000000 00000001 win32k!xxxRealSleepThread+0x1d7
        9a99fc60 95346414 000025ff 00000000 00000001 win32k!xxxSleepThread+0x2d
        9a99fcb8 95349966 9a99fce8 000025ff 00000000 win32k!xxxRealInternalGetMessage+0x4b2
        9a99fd1c 8283e1fa 000afb80 00000000 00000000 win32k!NtUserGetMessage+0x3f
        9a99fd1c 76f270b4 000afb80 00000000 00000000 nt!KiFastCallEntry+0x12a
        000afb3c 7705cde0 7705ce13 000afb80 00000000 ntdll!KiFastSystemCallRet
        000afb40 7705ce13 000afb80 00000000 00000000 USER32!NtUserGetMessage+0xc
        000afb5c 0055148a 000afb80 00000000 00000000 USER32!GetMessageW+0x33
        000afb9c 005516ec 00550000 00000000 0012237f notepad!WinMain+0xe6
...

```

The second way to switch process views in the host debugger is to perform an *invasive* process context switch on the target machine by using the */i* option of the *.process* command. This method

is particularly useful when you need to set breakpoints in user-mode code locations, given they're always interpreted relative to the current process context on the target machine, as you also learned back in Chapter 2. This method requires the target machine to exit the debugger break-in mode and run to complete the request.

After the target is let go by the host debugger, the kernel on that side thaws the frozen processors and exits the break-in loop. Before it does so, however, it also schedules a high-priority work item to transition over to the new process context that was requested by the host debugger.

```
1: kd> !process 0 0 notepad.exe
PROCESS 874fa030 SessionId: 1 Cid: 14ac Peb: 7ffdf000 ParentCid: 1348
Image: notepad.exe
1: kd> .process /i 874fa030
You need to continue execution (press 'g' <enter>) for the context
to be switched. When the debugger breaks in again, you will be in
the new process context.
1: kd> g
Break instruction exception - code 80000003 (first chance)
```

The work item that induced the previous debug break runs on a *leased* system thread that runs in the context of the requested process. The host debugger breaks right back in again before any of its threads have a chance to continue executing past where they were at the time of the original break-in. You can also confirm that the current thread context is a kernel thread, and not a thread from the user-mode process itself. Notice that thread is indeed owned by the system (kernel) *process*, which always has a PID value of 4, as reported by the *Cid* (client thread ID) you get from the *!thread* command.

```
0: kd> !thread
THREAD 856e94c0 Cid 0004.0038 Teb: 00000000 Win32Thread: 00000000 RUNNING on processor 0
ChildEBP RetAddr Args to Child
8a524c0c 82b30124 00000007 8293b2f0 856e94c0 nt!RtlpBreakWithStatusInstruction
8a524d00 8287da6b 00000000 00000000 856e94c0 nt!ExpDebuggerWorker+0x1fa
8a524d50 82a08fda 00000001 a158a474 00000000 nt!ExpWorkerThread+0x10d
8a524d90 828b11d9 8287d95e 00000001 00000000 nt!PspSystemThreadStartup+0x9e
00000000 00000000 00000000 00000000 00000000 nt!KiThreadStartup+0x19
```

Nevertheless, this system thread is attached to the target process you requested, which you can confirm using the *!process* kernel debugger extension command and *-I* to indicate you would like the current process context displayed.

```
0: kd> !process -I 0
PROCESS 874fa030 SessionId: 1 Cid: 14ac Peb: 7ffdf000 ParentCid: 1348
Image: notepad.exe
```

User-mode code breakpoints you enter in this host debugger break-in state will be resolved relative to this process context, exactly as desired.

Managed-Code Debugging

As previously mentioned in Chapter 2, one of the unfortunate limitations of the Windows debuggers is that they don't support source-level debugging of .NET applications. This doesn't mean that you can't use WinDbg to debug managed code; it simply means that you won't have the convenience of source-level debugging, such as single-stepping and source line breakpoints, when doing so. It's as if you were debugging system code without the source code; only it's worse because many important commands that work for native system debugging, such as displaying call stacks using the *k* command, don't even work for managed code. Fortunately, there is at least a workaround in the form of a WinDbg extension called *SOS*, which Microsoft ships with the .NET Framework. This useful extension is covered in more detail later in this section.

Because of this limitation, the Microsoft Visual Studio environment remains the debugger of choice for .NET debugging. To better understand why the Windows debuggers are lacking in this regard, it's useful to first discuss the architecture used by Visual Studio and the .NET Common Language Runtime (CLR) environment to implement their support for managed-code debugging and understand the way they collaborate to present a seamless native/managed debugging experience. Just like other .NET-related discussions in this book, the coverage centers on the architecture in version 4.0 of the .NET Framework.

Architecture Overview

The first challenge when designing an architecture that enables debugging of Microsoft Intermediate Language (MSIL) .NET code is that such code gets translated into machine instructions on the fly by the CLR's Just-in-Time (JIT) compiler. For performance reasons, this run-time code generation is done lazily only after a method is actually invoked. In particular, this means that to insert a code breakpoint, the debugger needs to wait until the code in question is loaded into memory so that it can edit the code in memory and insert the debug break instruction at the appropriate location. The native debug events generated by the OS aren't sufficient by themselves to support this type of MSIL debugging because only the CLR knows when the .NET methods are compiled or how the managed class objects are represented in memory.

For those reasons, the CLR designed an infrastructure for debuggers to inspect and control managed targets with the help of a dedicated thread that runs as part of every .NET process and has intimate knowledge of its internal CLR data structures. This thread is known as the *debugger runtime controller* thread, and it runs in a continuous loop waiting for messages from the debugger process. Even in the break-in state, the managed target process isn't entirely frozen because this thread must still run to service the debugger commands. Any .NET application will have this extra debugger thread even when it isn't being actively debugged with a managed-code debugger. To confirm this fact, you can use the following "Hello World!" C# sample from the companion source code.

```
C:\book\code\chapter_03\HelloWorld>test.exe
Hello World!
Press any key to continue...
```

You can now use the steps described in Chapter 2 to start a live kernel-debugging session and noninvasively observe the threads that the managed process contains when it's active. Notice the presence of the debugger runtime controller thread (the `clr!DebuggerRThread::ThreadProc` thread routine in the following listing) even though the .NET process isn't being debugged with a user-mode debugger.

```

tkd> .symfix
tkd> .reload
tkd> !process 0 0 test.exe
PROCESS 85520c88 SessionId: 1 Cid: 07b8 Peb: 7ffdf000 ParentCid: 0e5c
Image: test.exe
tkd> .process /r /p 85520c88
tkd> !process 85520c88 7
PROCESS 85520c88 SessionId: 1 Cid: 07b8 Peb: 7ffdf000 ParentCid: 0e5c
Image: test.exe
...
THREAD 9532ed20 Cid 07b8.1e5c ...
828d74b0 SynchronizationEvent
885f5cb8 SynchronizationEvent
86a0e808 SynchronizationEvent
...
0116f7fc 5d6bb4d8 00000003 0116f824 00000000 KERNEL32!WaitForMultipleObjects+0x18
0116f860 5d6bb416 d6ab8654 00000000 00000000 clr!DebuggerRThread::MainLoop+0xd9
0116f890 5d6bb351 d6ab8678 00000000 00000000 clr!DebuggerRThread::ThreadProc+0xca
0116f8bc 76f9ed6c 00000000 0116f908 779c377b clr!DebuggerRThread::ThreadProcStatic+0x83
0116f8c8 779c377b 00000000 6cb23a74 00000000 KERNEL32!BaseThreadInitThunk+0xe
0116f908 779c374e 5d6bb30c 00000000 00000000 ntdll!_RtlUserThreadStart+0x70
0116f920 00000000 5d6bb30c 00000000 00000000 ntdll!_RtlUserThreadStart+0x1b
tkd> q

```

Because of its reliance on this helper thread, the managed-code debugging paradigm is often referred to as *in-process* debugging, in contrast to the out-of-process debugging architecture used by native code user-mode debuggers, which requires no active collaboration from the target process. The contract defined by the CLR for managed-code debuggers to interact with the runtime controller thread is represented by a set of COM interfaces implemented in the `mscordbi.dll` .NET Framework DLL. Because this contract is published as a set of COM interfaces, you can write a managed-code debugger in C/C++, and also in any other .NET language, where the COM Interop facilities can be used to consume the CLR debugging objects implemented in this DLL.

The Visual Studio debugger is based on this same CLR debugging infrastructure, which it also uses to implement its support for managed-code debugging. The components used to service the user actions in the debugger are represented, at a high level, in Figure 3-6. The debugger front-end UI processes any commands entered by the user and forwards them to the debugger's back-end engine, which in turn internally uses the CLR debugging COM objects from `mscordbi.dll` to communicate with the runtime controller thread in the managed target process. These COM objects take care of all the

internal details related to the private interprocess communication channel between the debugger and target processes.

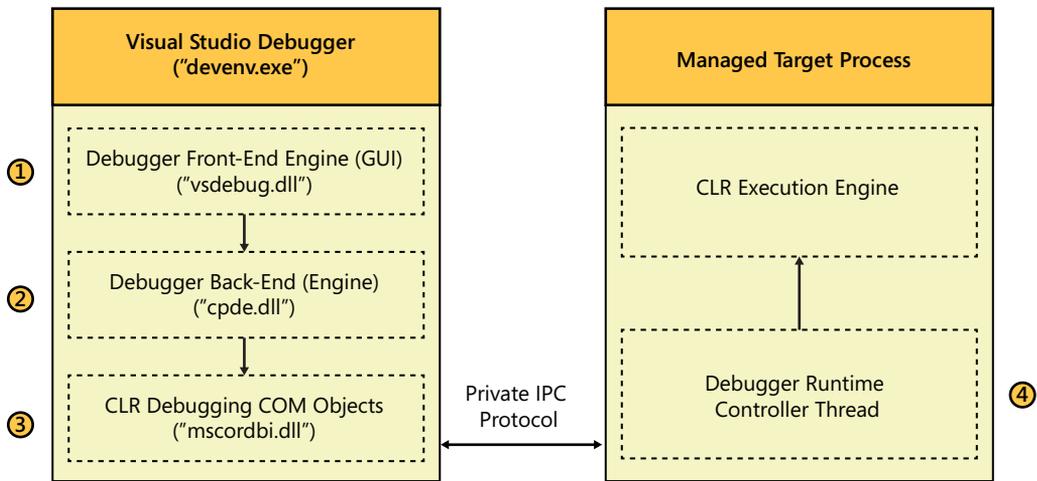


FIGURE 3-6 In-process managed debugging architecture in Visual Studio and the CLR.

This architecture has one big advantage, which is that it insulates the debuggers from the intricate details of the internal CLR execution engine data structures by having a higher-level contract and communication channel between the managed-code debuggers and the CLR debugger controller thread. This means the layouts of those data structures can change without breaking the functionality of those debuggers.

Unfortunately, this architecture also has several drawbacks. First, this model doesn't work for debugging of crash dumps because the target isn't running in that case, so the debuggers can't rely on an active debugger helper thread to perform their actions when debugging a memory crash dump file.

Second, the operating system is unaware that the application is being debugged using this private interprocess communication channel. Up until .NET version 4.0, Visual Studio debugging of managed applications didn't work at all on machines that also had a host kernel debugger attached to them. Because the OS didn't know that the managed process was being debugged, exceptions raised for the purpose of managed debugging were being incorrectly caught by the kernel debugger. The official workaround to this problem was documented in the Knowledge Base (KB) article at <http://support.microsoft.com/kb/303067>, but it's hardly satisfactory because it recommends disabling the kernel debugger entirely. Fortunately, this problem is now fixed in Visual Studio 2010—at least for managed applications compiled for .NET 4.0—because the debugger now also attaches to the target process debug port as a regular native user-mode debugger. However, the in-process managed-debugging architecture is still otherwise being used in that release as the main live, managed-code debugging channel.

Table 3-3 contains a comparative analysis of the in-process and out-of-process debugging architectures.

TABLE 3-3 In-Process and Out-of-Process Managed Debugging Paradigms

	Advantages	Drawbacks
In-process debugging	<ul style="list-style-type: none">■ Easy access to CLR data structures■ Faster single-stepping	<ul style="list-style-type: none">■ Poor integration with kernel-mode debugging■ Doesn't work for crash dump debugging
Out-of-process debugging	<ul style="list-style-type: none">■ Supports crash dump debugging■ Natural integration with native debugging■ No side effects preventing kernel-mode debugging	<ul style="list-style-type: none">■ More difficult for the debugger to stay in-sync with the CLR execution engine's data structures

Given the benefits of out-of-process debugging, the CLR and Visual Studio probably will continue to move toward that architecture for managed-code debugging in the future. That trend has already begun in .NET 4.0 and Visual Studio 2010, where the out-of-process architecture is now used to support crash dump debugging of managed processes.

The SOS Windows Debuggers Extension

Many WinDbg commands don't work natively when debugging a .NET target program. For instance, the *k* command cannot display the names of managed functions in a call stack and the *dv* command cannot display the values of local variables from those functions, either. To understand why, remember that MSIL images are compiled on the fly, so the dynamically generated code addresses are completely unknown to the symbols that the Windows debugger relies on to map the addresses to their friendly symbolic names. Even when an MSIL image is precompiled into a native one—a process known as *NGEN'ing* the assembly—the generated native image is actually machine-specific and won't have a corresponding symbol file, either. The .NET Framework DLL assemblies fall into this second category because they are usually NGEN'ed on the machines where they're installed to improve the performance of all the applications that use them.

How SOS Works

To work around the lack of native support for managed-code debugging in the Windows debuggers, the .NET Framework ships the *sos.dll* debugger extension module. This extension was doubly useful in earlier releases of the .NET Framework because it was also the only supported way to perform crash dump debugging of .NET code, given that Visual Studio started supporting out-of-process debugging of managed code only in its 2010 release.

This debugger extension is built as part of the CLR code base, so it has intimate knowledge of the internal layouts of the CLR data structures, allowing it to read the virtual address space of the target process directly and parse the CLR execution engine structures that it needs. These capabilities

enable it to support out-of-process managed-code debugging. When using SOS, you'll at least be able to display managed call stacks, set breakpoints in managed code, find the values of local variables, dump the arguments to method calls, and perform most of the inspection and control debugging actions that you can use in native-code debugging—only without the convenience of source-level debugging.

Symbols for .NET modules are used by managed-code debuggers only to enable source-level debugging (source lines, names of local function variables, and so on). Even without symbol files for managed assemblies, you still can do a lot of things you aren't able to do in native-code debugging, where the symbols are absolutely crucial. This is because MSIL images also carry metadata describing the type information for the classes they host, allowing any component with internal knowledge of how to parse that information to use it for displaying function names in a call stack, dump the values of local variables (though without their names), or find the parameters to function calls. This is precisely how the SOS Windows debugger extension enables out-of-process managed-code debugging—even without symbol files or any additional help from the CLR debugger runtime controller thread.

Debugging Your First .NET Program Using SOS

To provide a practical illustration for how to use SOS to debug .NET programs in WinDbg, you'll now use it to debug the following C# program from the companion source code, which you should compile to target CLR version 4.0, as described in the procedure provided in the Introduction of this book.

```
//  
// C:\book\code\chapter_03\HelloWorld>main.cs  
//  
public class Test  
{  
    public static void Main()  
    {  
        Console.WriteLine("Hello World!");  
        Console.WriteLine("Press any key to continue...");  
        Console.ReadLine();  
        Console.WriteLine("Exiting...");  
    }  
}
```

Every version of the CLR has its own copy of the SOS extension DLL that understands its internal data structures and is able to decode them. For this reason, you must always load the version of the extension that comes with the CLR version that's used by the target process you're trying to debug. In addition, the SOS commands work only after the CLR execution engine DLL has been loaded, so you need to wait for its module load event to occur. This happens early during the startup of the .NET target as the CLR shim DLL (*mscoree.dll*) hands the reins over to the CLR execution engine DLL, which is *clr.dll* in the case of CLR version 4 (.NET 4.x), and *mscorwks.dll* in the case of CLR version 2

(.NET 2.x and .NET 3.x). You can get notified of this module load event in the debugger by using the `sxe ld` command, as shown in the following listing.

```
0:000> vercommand
command line: '"c:\Program Files\Debugging Tools for Windows (x86)\windbg.exe"
c:\book\code\chapter_03\HelloWorld\test.exe'
0:000> .symfix
0:000> .reload
0:000> sxe ld clr.dll
0:000> g
ModLoad: 5fad0000 6013e000 C:\Windows\Microsoft.NET\Framework\v4.0.30319\clr.dll
ntdll!KiFastSystemCallRet:
779970b4 c3 ret
0:000> .lastevent
Last event: 1e30.c20: Load module C:\Windows\Microsoft.NET\Framework\v4.0.30319\clr.dll at
5fad0000
```

After the execution engine DLL is loaded, you can load the SOS extension module before any managed code has a chance to run inside the target process. A command you'll find useful when loading the SOS extension DLL is the `.loadby` debugger command. This command works just like the more basic `.load` command, but it looks up the extension module under the same path where its second module parameter was loaded from. By specifying the CLR execution engine DLL module name, you will be sure to load the `sos.dll` extension from the same location so that it matches the precise CLR version of the target. One of the useful SOS commands is the `!eeversion` command, which displays the current version of the CLR in the target process.

```
0:000> .loadby sos clr
0:000> !eeversion
4.0.30319.239 retail
0:000> g
```

The program now waits for user input in the `ReadLine` method. If you break into the debugger at this point by using the `Debug\Break` menu action, you'll see that the `k` command isn't able to properly display the function names in the managed code frames from the main thread in the .NET process. (Notice the very large offsets in the frames from the `mscorlib_ni` native image of the `mscorlib.dll` .NET Framework assembly, which is indicative of missing or unresolved symbols.) The unmanaged frames are still decoded correctly.

```
0:004> ~0s
0:000> k
ChildEBP RetAddr
0017e998 77996464 ntdll!KiFastSystemCallRet
0017e99c 75ea4b6e ntdll!ZwRequestWaitReplyPort+0xc
0017e9bc 75eb2833 KERNEL32!ConsoleClientCallServer+0x88
0017eab8 75efc978 KERNEL32!ReadConsoleInternal+0x1ac
0017eb40 75ebb974 KERNEL32!ReadConsoleA+0x40
0017eb88 5efc1c8b KERNEL32!ReadFileImplementation+0x75
0017ec08 5f637cc8 mscorlib_ni+0x2c1c8b
0017ec30 5f637f60 mscorlib_ni+0x937cc8
0017ec58 5ef78bfb mscorlib_ni+0x937f60
0017ec74 5ef5560a mscorlib_ni+0x278bfb
0017ec94 5f63e6f5 mscorlib_ni+0x25560a
```

```

0017eca4 5f52a7aa mscorlib_ni+0x93e6f5
0017ecb4 5fad21bb mscorlib_ni+0x82a7aa
0017ecc4 5faf4be2 clr!CallDescrWorker+0x33
0017ed40 5faf4d84 clr!CallDescrWorkerWithHandler+0x8e
0017ee7c 5faf4db9 clr!MethodDesc::CallDescr+0x194
0017ee98 5faf4dd9 clr!MethodDesc::CallTargetWorker+0x21
0017eeb0 5fc273c2 clr!MethodDescCallSite::Call_ReturnSlot+0x1c
0017f014 5fc274d0 clr!ClassLoader::RunMain+0x24c
0017f27c 5fc272e4 clr!Assembly::ExecuteMainMethod+0xc1
0017f760 5fc276d9 clr!SystemDomain::ExecuteMainMethod+0x4ec
0017f7b4 5fc275da clr!ExecuteEXE+0x58
...

```

Fortunately, the `!clrstack` command from the SOS debugger extension allows you to see the managed frames in the thread's call stack.

```

0:000> !clrstack
OS Thread Id: 0xe48 (0)
Child SP IP          Call Site
0017eba8 779970b4 [InlinedCallFrame: 0017eba8]
0017eba4 5efc1c8b DomainNeutralILStubClass.IL_STUB_PInvoke(Microsoft.Win32.SafeHandles.
SafeFileHandle, Byte*, Int32, Int32 ByRef, IntPtr)
0017eba8 5f637cc8 [InlinedCallFrame: 0017eba8] System.IO.__ConsoleStream.ReadFile(Microsoft.
Win32.SafeHandles.SafeFileHandle, Byte*, Int32, Int32 ByRef, IntPtr)
0017ec1c 5f637cc8 System.IO.__ConsoleStream.ReadFileNative(Microsoft.Win32.SafeHandles.
SafeFileHandle, Byte[], Int32, Int32, Int32, Int32 ByRef)
0017ec48 5f637f60 System.IO.__ConsoleStream.Read(Byte[], Int32, Int32)
0017ec68 5ef78bfb System.IO.StreamReader.ReadBuffer()
0017ec7c 5ef5560a System.IO.StreamReader.ReadLine()
0017ec9c 5f63e6f5 System.IO.TextReader+SyncTextReader.ReadLine()
0017ecac 5f52a7aa System.Console.ReadLine()
0017ecb4 0043009f Test.Main() [c:\book\code\chapter_03\HelloWorld\main.cs @ 9]
0017eee4 5fad21bb [GCFrame: 0017eee4]

```

The `mscorlib.ni.dll` module shown in the stack trace output of the `k` command is the NGEN image ("ni") corresponding to the `mscorlib.dll` MSIL image. You can treat these modules just like their MSIL sources for the purpose of SOS debugging. In particular, you can set breakpoints at managed code functions from both MSIL or NGEN images by using the `!bpmd` SOS extension command.

For example, you can set a breakpoint at the `WriteLine` method that would be executed by the next line of source code. This .NET method is defined in the `System.Console` class of the `mscorlib.dll` .NET assembly (or in this case, its `mscorlib.ni.dll` NGEN version). The `!bpmd` command takes the target module name as its first argument (without the extension!) and the fully qualified name of the .NET method as its second argument, as shown in the following listing.

```

0:004> !bpmd mscorlib_ni System.Console.WriteLine
Found 19 methods in module 5ed01000...
MethodDesc = 5ed885a4
Setting breakpoint: bp 5EFAD4FC [System.Console.WriteLine()]
MethodDesc = 5ed885b0
Setting breakpoint: bp 5F52A770 [System.Console.WriteLine(Boo!ean)]
MethodDesc = 5ed885bc
...

```

Adding pending breakpoints...

0:004> g

This command adds breakpoints to all overloads of the *WriteLine* method (19 of them in the previous case). If you now press Enter in the active command prompt window from the target process, you'll notice that the debugger hits your breakpoint next.

```
Breakpoint 13 hit
mscorlib_ni+0x2570ac:
5ef570ac 55          push     ebp
```

You can again use the *!clrstack* command to see the current stack trace at the time of this breakpoint. The *-a* option of this command also allows you to view the arguments to the managed frames on the stack.

```
0:000> !clrstack -a
OS Thread Id: 0x18e0 (0)
Child SP IP      Call Site
0018f260 5ef570ac System.Console.WriteLine(System.String)
PARAMETERS:
    value (<CLR reg>) = 0x01fdb24c
0018f264 004600ab Test.Main()
*** WARNING: Unable to verify checksum for test.exe
    [c:\book\code\chapter_03\HelloWorld\main.cs @ 10]
0018f490 5fad21bb [GCFrame: 0018f490]
```

Notice how this command also displays the address of the .NET string object that was passed to the *WriteLine* method, which you can dump using the *!do* ("dump object") SOS debugger extension command.

```
0:000> !do 0x01fdb24c
Name:      System.String
MethodTable: 5f01f92c
EEClass:   5ed58ba0
Size:      34(0x22) bytes
String:    Exiting...
Fields:
    MT      Field  Offset          Type VT      Attr      Value Name
5f0228f8 4000103      4      System.Int32  1 instance      10 m_stringLength
5f021d48 4000104      8      System.Char   1 instance      45 m_firstChar
5f01f92c 4000105      8      System.String 0 shared      static Empty
>> Domain:Value 002a1270:01fd1228 <<
```

Notice that the *!clrstack* command doesn't display the unmanaged functions on the call stack, though it's usually easy to see where the managed calls fit in the overall stack trace by combining the *!clrstack* and the regular *k* back-trace command, which should give you everything you need to know about what code the current thread is currently executing. Note that SOS also has a *!dumpstack* command that attempts to do this merge, but its output can be rather noisy.

The SOS extension also has several other useful commands that you can use to inspect .NET programs, including a variant of the *u* ("un-assemble") command that's also able to decode the addresses of managed function calls in addition to unmanaged addresses. For example, you could use

this command to obtain the disassembly of the current function at the time of the breakpoint in the previous case (the *WriteLine* method).

```
0:000> !u .
preJIT generated code
System.Console.WriteLine(System.String)
Begin 5ef570ac, size 1a
>>> 5ef570ac 55          push   ebp
5ef570ad 8bec          mov    ebp,esp
5ef570af 56          push   esi
5ef570b0 8bf1          mov    esi,ecx
5ef570b2 e819000000  call   mscorlib_ni+0x2570d0 (5ef570d0) (System.Console.get_Out()),
mdToken: 060008fd
...
```

Notice how the regular *u* command, by contrast, doesn't display the friendly name of the function itself or of the call to *get_Out* (a managed method too) that's made inside the same function.

```
0:000> u .
mscorlib_ni+0x2570ac:
5ef570ac 55          push   ebp
5ef570ad 8bec          mov    ebp,esp
5ef570af 56          push   esi
5ef570b0 8bf1          mov    esi,ecx
5ef570b2 e819000000  call   mscorlib_ni+0x2570d0 (5ef570d0)
```

If you would like to experiment with more SOS debugger commands, you can find a listing of those commands and a brief summary of what they do by using the *!help* command in the WinDbg debugger.

```
0:000> !help
-----
SOS is a debugger extension DLL designed to aid in the debugging of managed
programs. Functions are listed by category, then roughly in order of
importance. Shortcut names for popular functions are listed in parenthesis.
Type "!help <functionname>" for detailed info on that function.
...
0:000> $ Terminate this debugging session now...
0:000> q
```

Table 3-4 recaps the basic SOS commands introduced during this experiment.

TABLE 3-4 Basic SOS Extension Commands

Command	Purpose
<i>!eeversion</i>	Display the target CLR (execution engine) version.
<i>!bpmd</i>	Set a breakpoint using a managed .NET method.
<i>!do</i> (or <i>!dumpobj</i>)	Dump the fields of a managed object.
<i>!clrstack</i> <i>!clrstack -a</i>	Display the managed frames in the current thread's call stack. The optional <i>-a</i> option is used to also display the arguments to the functions on the call stack. These values are the extension's best guess, however; so, they're not always accurate.
<i>!u</i>	Display the disassembly of a managed function.

Despite the fact you can achieve a lot of critical debugging tasks using the SOS extension, the managed-code debugging experience in the Windows debuggers still leaves a lot to be desired. The Windows debuggers clearly are not your first choice when debugging the managed code you write yourself, but SOS can still be a good option, especially if you can't get Visual Studio installed on the target machine or if you are debugging without source code—in which case, you don't lose much by using WinDbg anyway.

Script Debugging

Visual Studio also supports source-level debugging of script languages such as VBScript or JScript. This is internally implemented using the same in-process paradigm that managed-code debugging relies on. One of the reasons that the CLR used the in-process debugging model when it was first released to the public in 2002 was that script debugging had been successfully using it for years (since the mid-90s). In both cases, the debugger needs the script host's or the CLR execution engine's collaboration to support source-level debugging of the target process.

Architecture Overview

To understand how script debugging works, it's useful to first explain a few basic concepts about how script languages are executed in Windows. The key to that architecture is the *Active Scripting* specification. This specification was introduced by Microsoft in the 90s and defines a set of COM interfaces to allow script languages that implement them to be hosted in any conforming host application. In Windows, both VBScript and JScript (Microsoft's implementation of JavaScript) are Active Scripting languages whose implementation fully conforms with that specification.

The Active Scripting specification defines a language-processing *engine*, with the Active Scripting *host* using that engine when the script needs to be interpreted. Examples of Active Scripting engines are *vbscript.dll* and *jscript.dll*, which both ship with Windows under the *system32* directory. Examples of Active Scripting hosts include the Internet Information Services (IIS) web server (server-side scripts embedded in ASP or ASP.NET pages), Internet Explorer (client-side script hosting in web pages), and the Windows scripting hosts (*cscript.exe* or *wscript.exe*) that ship with Windows and can be used to host scripts executed from a command prompt. There are also third-party Active Scripting engines to support other script languages, including Perl and Python.

In addition, the Active Scripting specification also defines a contract (a set of COM interfaces, again) for debuggers to take advantage of the host in their operations. An Active Scripting host that supports debugging (that is, implements the required COM interfaces) is called a *smart* host. All recent versions of Internet Explorer, IIS, and the Windows scripting hosts are smart hosts that implement those interfaces, which is at the heart of the magic that enables Visual Studio to debug scripts hosted by any of those processes. A *Process Debug Manager* (PDM) component (*pdm.dll*) is shipped with the Visual Studio debugger to insulate script engines from having to understand the intricacies of script debugging. In many ways, the PDM component serves the same purpose that the CLR runtime debugger controller thread and *mscordbi.dll* serve during managed debugging, as illustrated in Figure 3-7.

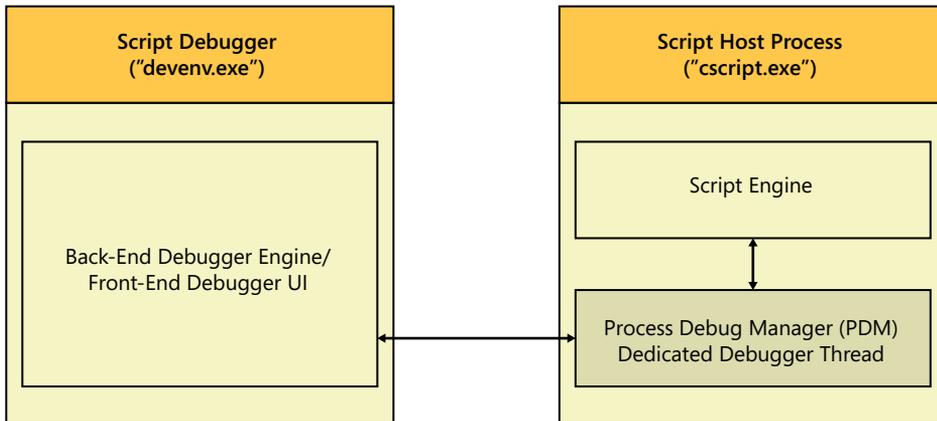


FIGURE 3-7 In-process script debugging architecture.

One way that Active Scripting debugging differs from managed-code debugging is that smart hosts usually do not expose their debugging services by default, whereas in the case of the CLR there is always a debugger thread running in the managed-code process. In Internet Explorer, for instance, you need to first enable script debugging in the host process by clearing the Disable Script Debugging option on the Tools\Internet Options\Advanced tab, as shown in Figure 3-8.

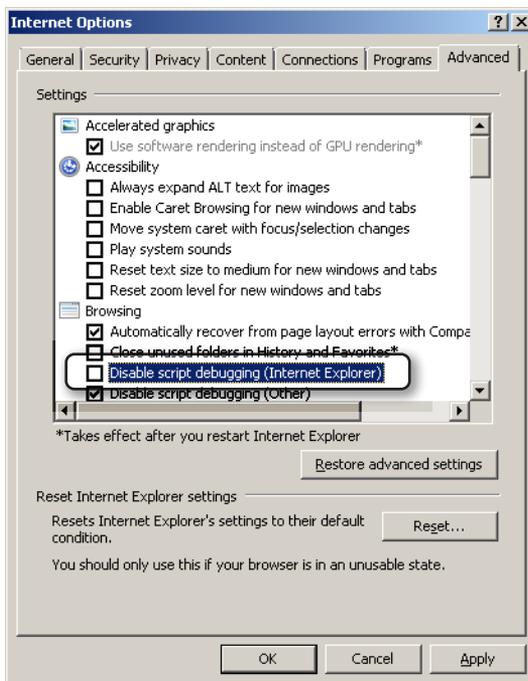


FIGURE 3-8 Enabling Internet Explorer script debugging.

In the same way, you need to explicitly enable debugging of scripts in the Windows scripting hosts (*cscript.exe/wscript.exe*) using the *//X* option if you want to also debug the executed scripts. The IIS web server manager also has a UI option for enabling server-side script debugging.

Debugging Scripts in Visual Studio

The following sample script from the companion source code, which simply displays its start time and arguments list to the console before exiting, will serve as a good example for how to use the *//X* option to enable script-debugging support in the *cscript.exe* host process and step through console scripts in Visual Studio.

```
//  
// C:\book\code\chapter_03\script>test.js  
//  
var g_argv = new Array();  
  
//  
// store command-line parameters and call the main function  
//  
for (var i=0; i < WScript.Arguments.length; i++)  
{  
    g_argv.push(WScript.Arguments(i));  
}  
  
WScript.Quit(main(g_argv));  
  
function main(argv)  
{  
    WScript.Echo("Script started at " + new Date());  
    if (g_argv.length > 0)  
    {  
        WScript.Echo("Arguments: " + g_argv);  
    }  
}
```

Notice in particular the double slash in the *//X* option in the following command, which *cscript.exe* and *wscript.exe* use to distinguish their own options from the executed script's options.

```
C:\Windows\system32\cscript.exe //X C:\book\code\chapter_03\script\test.js 1 2  
Microsoft (R) Windows Script Host Version 5.8  
Copyright (C) Microsoft Corporation. All rights reserved.  
...
```

When you run the preceding command on a machine with Visual Studio 2010 installed, you'll be presented with a Visual Studio "attach" dialog box similar to the one shown in Figure 3-9. You might get another dialog box to consent to UAC elevation if you invoked the script from an elevated administrative command prompt, given that the Visual Studio debugger also needs to run elevated in that case.

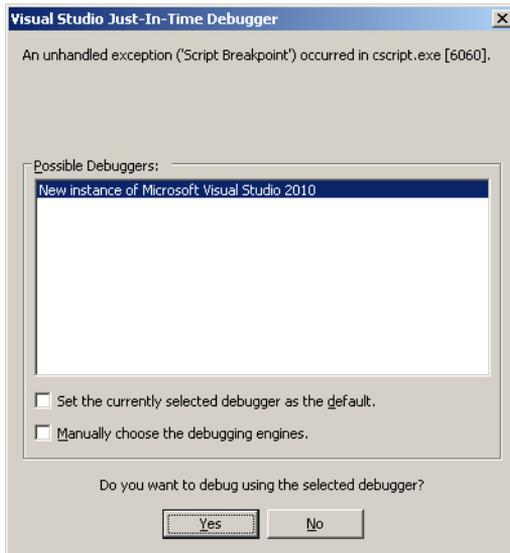


FIGURE 3-9 Attaching to a script using the Visual Studio debugger.

You are then able to use the Visual Studio debugger to step through (using the F10 and F11 shortcuts) the script and debug it (with source-level information!), just as you would with any native-code or managed-code application, as shown in Figure 3-10.

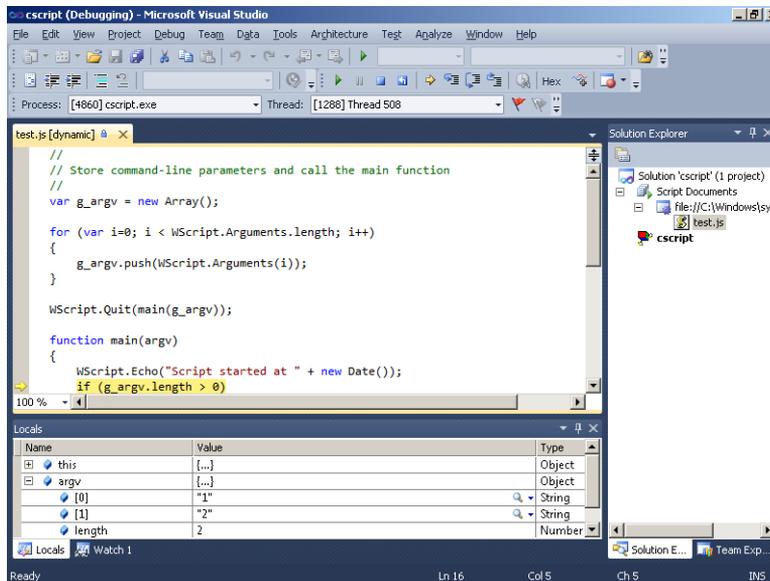


FIGURE 3-10 Source-level script debugging using Visual Studio 2010.

Remote Debugging

Remote debugging is a convenient feature that allows you to control a target remotely using a debugger instance that's running on a different machine on the network. This section provides an inside look at how this feature is implemented by debugger programs in Windows, as well as how you can use it in the case of the WinDbg and Visual Studio debuggers.

Architecture Overview

At a high level, two conceptual models are used to support remote debugging: remote *sessions* and remote *stubs*. In both cases, a process needs to be running on the same machine as the target so that it can carry out the commands that are typed in the remote debugger. In the case of a remote session, the debugger session is entirely on the target machine and the remote debugger instance simply acts as a "dumb" client to send commands to the local debugger instance. In the case of a remote stub, the debugger session is running remotely, with a "stub" broker process running locally on the target machine and acting as a gateway to get information in and out of the target.

Remote sessions are used when collaboration among multiple engineers is needed to investigate a certain failure. In that case, a local debugger instance runs on the repro machine, and developers are then able to start typing commands remotely from their respective machines, take a look at the failure, and even leave comments and see each other's debugging attempts and commands as they get typed in. WinDbg supports this very useful form of remote debugging, which is illustrated in Figure 3-11.

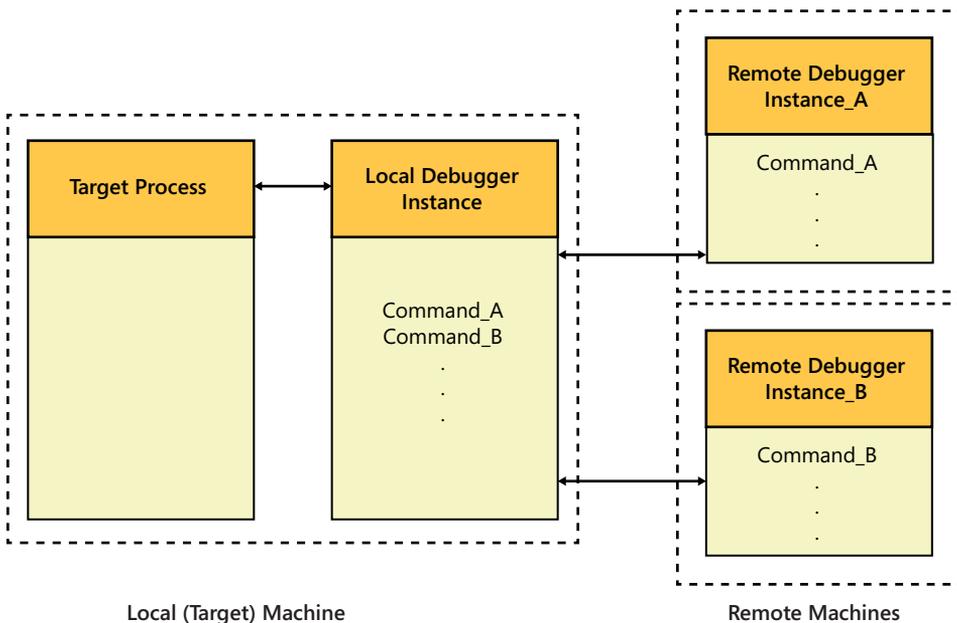


FIGURE 3-11 Remote debugging using remote sessions.

Remote stubs are used when it's important to set up the debugging environment on a remote computer, either because the full debugging environment can't be installed or because the symbols and sources cannot be accessed directly on the target computer. Both the Windows and Visual Studio debuggers support this form of remote debugging, though this architecture is more useful in the case of the Visual Studio debugger. (In fact, remote stubs are the only type of remote debugging that's supported by Visual Studio.) This is because a full Visual Studio debugging environment installation on the target machine is heavy handed, both in terms of the disk space it requires and the time it takes to complete, often making it an inadequate option for production environments. By comparison, the Visual Studio remote-debugging component, which also includes the remote stub process, is more lightweight and can be installed much faster when you need it on the target machine. Figure 3-12 illustrates this second form of remote debugging.

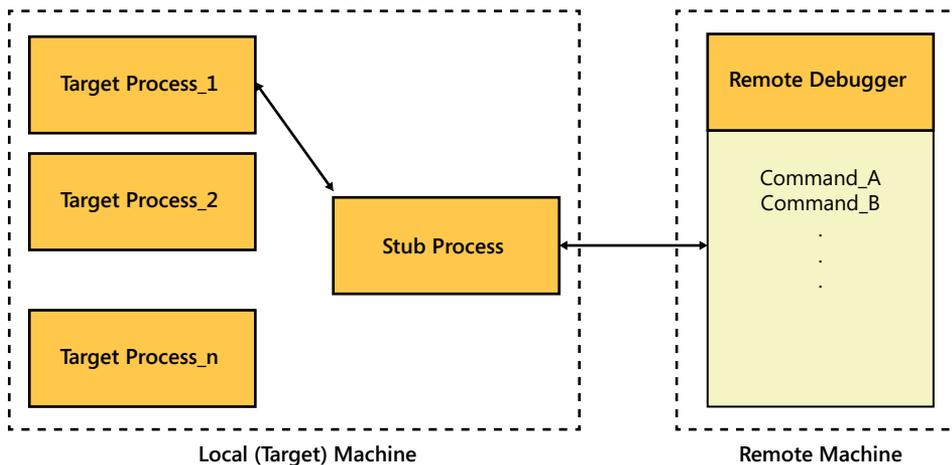


FIGURE 3-12 Remote debugging using remote stubs.

Remote Debugging in WinDbg

As mentioned earlier, WinDbg supports both remote sessions and remote stubs. This section will walk you through the steps for making use of either setup in your remote WinDbg debugging experiments.

Remote Sessions

Starting a WinDbg remote debugging session is straightforward. You use the `.server` debugger command, as illustrated by the following procedure.

Using Remote Sessions in WinDbg

1. From the local WinDbg instance controlling the target process, start a TCP/IP remote session by using the `.server` command, as illustrated in the following listing. This example uses port 4445, but any TCP port that's not currently in use would also work. Using the `.server` command with no arguments reminds you of its syntax.

```
0:000> vercommand
command line: '"c:\Program Files\Debugging Tools for Windows (x86)\windbg.exe"
notepad.exe'
0:000> $ .server without any arguments displays the usage...
0:000> .server
Usage: .server tcp:port=<Socket> OR .server npipe:pipe=<PipeName>
0:000> $ Open a remote debugging server session using port 4445
0:000> .server tcp:port=4445
Server started. Client can connect with any of these command lines
0: <debugger> -remote tcp:Port=4445,Server=TARIKS-LP03
```

2. On the remote machine where you plan to control the target, start a new WinDbg instance (the remote debugger) and connect to the previous TCP port remotely by using the `-remote` command-line option. You can also use the connection string provided in the output from the `.server` command in the preceding listing to remind you of the syntax. Note that the remote and target machines can be the same, so you can also execute this step on the same machine as step 1.

```
windbg.exe -remote tcp:Port=4445,Server=TARIKS-LP03
```

Notice how any commands you type in the new WinDbg instance's command window (the "remote" debugger) also appear in the first WinDbg instance on the target machine (the "local" debugger). The remote debugger is essentially acting only as a terminal for typing commands that get transferred and, ultimately, processed by the local debugger instance on the target machine.

3. You can terminate the entire debugging session (both the remote and local debugger instances) from the remote machine using the `qq` command. By contrast, the `q` command terminates only the remote instance and leaves the local debugger instance intact.

```
0:000> $ Remote debugger command prompt
0:000> qq
```

In addition to their obvious benefits in remote debugging scenarios, WinDbg remote sessions can be useful even when debugging on the same machine. An example is when you start debugging using one of the command-line Windows debuggers (`cdb.exe` or `ntsd.exe`) and later decide to switch to using WinDbg as a front-end UI to that same session.

Remote Stubs

Commands typed in the remote debugger command window in a remote debugging session are executed as if they were typed on the target machine. In particular, this means that debugger extensions are also loaded into the debugger process on the target machine and that symbols need to be accessible from that machine too. Although remote sessions are by far the more common scenario in WinDbg remote debugging, remote stubs can be used if you don't want to copy the symbols or otherwise make them available over the network so that they're accessible from the target machine.

You can use remote stubs in both user-mode and kernel-mode remote debugging with the Windows debuggers. The *dbgsvr.exe* process, which comes as part of the Windows debuggers package, is used as a stub process in remote user-mode debugging. The *kdsrv.exe* process, also included in the Windows debuggers package, is used in remote kernel-mode debugging.

Just as in the remote session case, you can follow the procedure described here by also running the remote debugger on the target machine, in case you don't have two separate machines.

Using Remote Stubs in WinDbg

1. Open a TCP/IP communication port on the target machine using the *dbgsvr.exe* stub process. This command displays a dialog box when it fails, but it won't show any messages on success.

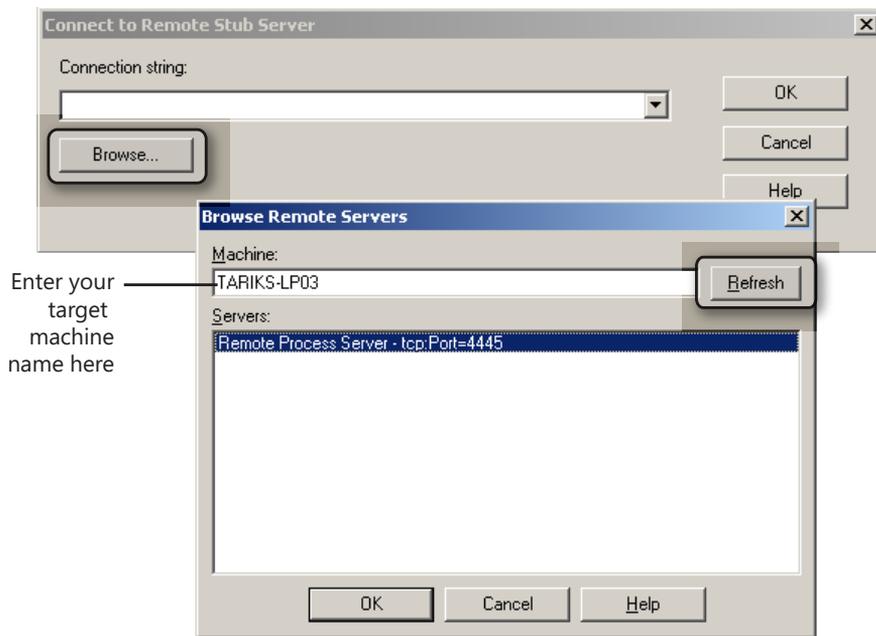
```
C:\Program Files\Debugging Tools for Windows (x86)\dbgsvr.exe -t tcp:port=4445
```

The *dbgsvr.exe* stub process is running in the background at this point. Using the *netstat.exe* tool that comes with Windows under the *system32* directory, you can display the open network ports on the machine and confirm that this stub process is listening for connections from a remote debugger on TCP port 4445:

```
C:\windows\system32\netstat.exe -a
Active Connections
  Proto Local Address           Foreign Address         State
  ...
  TCP   0.0.0.0:4445            TARIKS-LP03:0         LISTENING
  ...
```

2. On the remote machine, start a new *windbg.exe* instance and connect to the stub process on the target machine using the File\Connect to a Remote Stub menu action. Leave the Connection String combo box empty, and enter the target machine name in the new Browse

dialog box. WinDbg then automatically enumerates all the open stub sessions on the target machine for you, as demonstrated in the following screen shot:



3. After you connect to the remote stub, you can attach to processes running on the target machine by using the familiar File\Attach To A Process menu command (the F6 shortcut). However, this option now shows processes from the target machine, which is exactly what you want in this case. Once attached to a process, you can debug it as if it were running locally, with symbols and debugger extensions located relative to the same remote debugger machine (and not the target machine, as was the case in remote sessions).
4. When you no longer need the remote debugging channel, make sure you terminate the stub process on the target machine so that you release the TCP port it opened earlier for other uses on the machine. You can do that in the Windows task manager UI or by using the *kill.exe* command-line utility from the Windows debuggers package.

```
C:\Program Files\Debugging Tools for Windows (x86)>kill.exe dbgsrv
process dbgsrv.exe (4188) - '' killed
```

This same approach can be used for remote kernel-mode debugging, except you should use the *kdsrv.exe* stub instead of the *dbgsrv.exe* stub. Note that in that case, there are actually three machines involved: the regular target and host kernel debugger machines, and the remote machine you are using to run the debugger instance. The *kdsrv.exe* process is started as a remote stub on the kernel

host debugger machine, not the target machine that is being debugged with the kernel debugger. Symbols and extensions are again resolved relative to the remote debugger machine.

Remote Debugging in Visual Studio

Unfortunately, Visual Studio does not support the concept of remote debugging sessions or remote connection strings, which makes it difficult to use it for sharing a debugging session with another developer at the exact point of a particular failure in the way that WinDbg's `.server` command works. Instead, Visual Studio remote debugging uses the remote stub idea, with the `msvsmon.exe` process acting as the remote stub process to which the Visual Studio debugger process (`devenv.exe`) then connects from the remote machine. MSVSMON is also sometimes called the Visual Studio *debug monitor*, highlighting the fact that it controls the execution of the target on the local machine on behalf of the Visual Studio debugger on the remote machine.

The default communication protocol between the remote and the target machines uses Distributed COM (DCOM) with Windows authentication, so you need to configure the user running the debugger with correct security permissions on the target machine. If the account you're using is a domain user, for example, it will need to be an administrator on the target machine or a member of the *Debugger Users* security group created by Visual Studio.

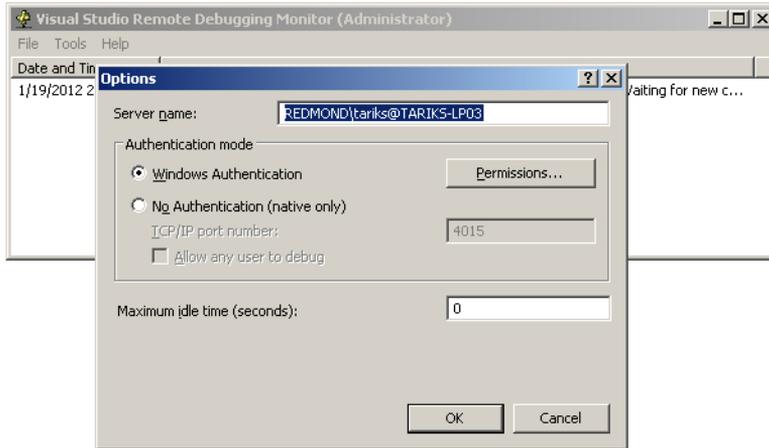
Using Remote Debugging in Visual Studio

1. Start the following C# sample program from the companion source code (and leave it waiting for user input) on the target machine.

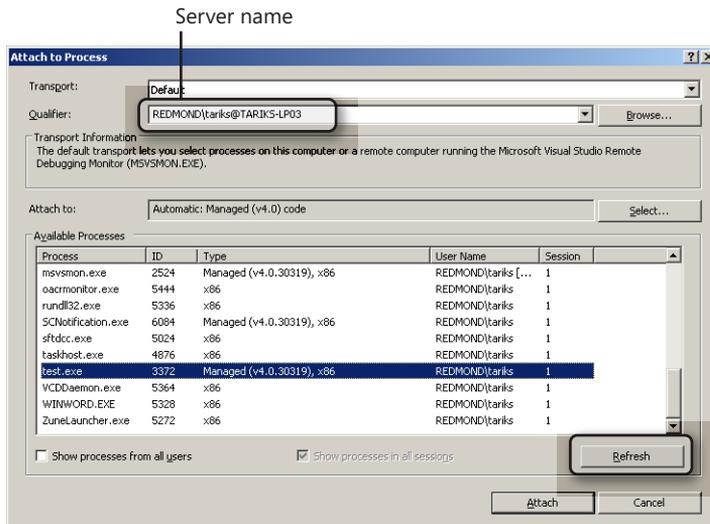
```
C:\book\code\chapter_03\HelloWorld>test.exe
Hello World!
Press any key to continue...
```

2. Install the Visual Studio 2010 remote-debugging components on the target machine from <http://www.microsoft.com/download/en/details.aspx?id=475>. This setup installation shouldn't take too long (at least compared to the full Visual Studio installation!) because the download is only a few megabytes large. Cancel the configuration wizard that comes up at the end of the installation.
3. Start the `msvsmon.exe` stub process on the target machine. Use the Tools\Options menu action of MSVSMON to change the server name for the connection, or simply leave the default value unchanged, as shown in the following screen shot.

```
C:\Program Files\Microsoft Visual Studio 10.0\Common7\IDE\Remote Debugger\x86>msvsmon.exe
```



- On the remote machine, which has the full Visual Studio development environment, use the Tools\Attach To Process menu action to attach to the *test.exe* process remotely. In the Qualifier text box, specify the server name you chose in step 3.



- If your Windows firewall is enabled (the default behavior), a consent dialog box appears on the target machine. After authorizing the firewall exception, you'll see a list of all the processes that are currently running on the target machine, and you can then attach to the managed *test.exe* process.

You might also see the warning shown in Figure 3-13 as Visual Studio tries to locate the symbols for the managed process. This is because managed-code symbols are resolved relative to the target machine.

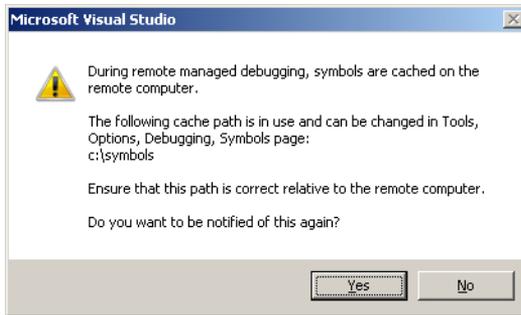


FIGURE 3-13 Managed-code symbols path resolution.

Unlike native code symbols, which are resolved in Visual Studio remote debugging relative to the remote machine (as was the case in the WinDbg stub-based remote debugging case), managed-code symbols are resolved relative to the target machine because of the in-process nature of managed-code debugging. Keep this in mind when you use Visual Studio and you want to perform remote source-level debugging of .NET applications because you need to copy the symbols to the target machine for the debugger to locate them successfully.

Summary

This chapter explained several debugging types and mechanisms in Windows and how they work at the system level. After reading this chapter, the following points should be familiar to you:

- Native user-mode debugging relies on a system-provided architecture that's based on a set of debug events generated by the system on behalf of the target process. These events are sent to a shared debug port object that the OS executive associates with the target process, with a dedicated thread in the debugger process waiting for new events and processing them in that order. This generic framework provides the foundation that user-mode debuggers use to control the target's execution, including setting code breakpoints and single-stepping the target.
- Kernel-mode debugging shares many concepts with user-mode debugging. Because of its global scope, however, a few debugging actions—such as setting code breakpoints and single-stepping the target—are implemented in a slightly different way to accommodate this reality.

- Both script and managed-code debugging use the in-process debugging paradigm, which has a helper thread running in the target process to assist the debugger in its operations. Visual Studio uses this architecture to support both types of debugging.
- The SOS Windows debugger extension can also be used to debug managed code in a completely out-of-process way, though without the convenience of source-level debugging.
- The Windows debuggers support collaborative remote user-mode and kernel-mode debugging using remote sessions. Visual Studio also supports remote debugging, but it uses a different architecture that takes advantage of a remote stub process running on the target machine. The type of remote debugging architecture also affects where symbols and debugger extensions get loaded from, so they each have their practical applications.

Now that you've seen how debuggers can actively debug local and remote targets in Windows, the next chapter will introduce you to another debugging approach when direct live debugging isn't an option. This very important debugging concept is called *postmortem* debugging.

Index

Symbols and Numbers

- \$\$>< command, 177
- \$ (dollar) sign, 40, 178
- 0xeb byte, 265
- 32-bit Windows, incomplete call stacks, 439
- 64-bit processors, 5
- 64-bit Windows, 183
 - incomplete call stacks, 440–441
 - Kernel Patch Protection, 381
 - stack trace capabilities, 432
- 1394 (FireWire) cable, 69
- 1394 IEEE ports
 - kernel debugging over, 69–73
 - types, 69
- @ (at) sign, 178
- @! character sequence, 180–181
- ?? command, 268, 278
- ~ command, 46, 254, 346, 511, 519
 - s suffix, 46–47, 511
- . (dot) alias, 263
- ... (ellipsis), 41
- * (wildcard character)
 - in bm command, 510
 - in k command, 512
 - in x command, 509

A

- A2W ASCII-to-Unicode ATL conversion macro, 298, 300
- access-check failures
 - debugging, 352–362
 - troubleshooting, 443
- access checks, 353
 - initiation of, 354
 - integrity levels and, 358–360
 - special user privileges and, 362

- UAC, 360–362
- access control entries (ACEs). *See* ACEs (access control entries)
- access control lists (ACLs), 353
- access denied failures
 - breaking on, 250
 - call site, finding, 245–246
 - observing, 247–248
- access tokens, 10, 353, 354
 - displaying, 355–357
 - dumping, 360–361
 - effective, 354, 362
- access violations
 - !analyze command, 269–270
 - debugging, 267–270
 - heap corruptions, 271
 - invalid memory, 198
 - page-heap fill patterns as pointers, 278
 - during process rundown, 339
 - unloaded memory access attempts, 340–343
- ACEs (access control entries), 353
 - mandatory integrity labels, 359, 361–362
- ACLs (access control lists), 353
- a command, 263
- Action\Snapshots menu action, 76
- Active Scripting, 112–114
- Active Template Library (ATL), 22, 29
- ActiveX, 20
- administrative privileges
 - for debugging, 87
 - default rights, 360
- advanced local procedure calls (ALPCs). *See* ALPCs (advanced local procedure calls)
- advapi32\CryptEncrypt API, 248
- advapi32.dll, 18, 226
- advapi32\EnumerateTraceGuidsEx API, 428
- advapi32\EventRegister API, 428

advapi32!EventWrite API

- advapi32!EventWrite API, 428
- advapi32!LogonUserW API, handling leaks from, 300–302
- AeDebug registry key
 - Auto value, 128
 - Debugger value, 139
 - populating, 127
 - Visual Studio settings, 132
- allocations, 480. *See also* memory allocations
- allocators, 271. *See also* memory allocations
 - NT heap functions, redirecting to, 272
- ALPC flag, 426
- ALPCs (advanced local procedure calls), 15–16
 - in console applications, 371–373
 - console UI processing, 368
 - message structure, dumping, 379
- Alpha processors, 4
- AMD-64, 4
- AMD-V processors, 73
- !analyze command, 143–144, 153, 269–270
 - v option, 269
- antivirus products, System Service Dispatch Table use, 381
- APCs (asynchronous procedure calls), 15
- API calls, 365
 - kernel-mode side, 385–386
 - transition to kernel mode, 383–384
 - user-mode side, 381–383
- APIs. *See also* Win32 APIs
 - asynchronous processing, showing, 342
 - documentation of, 16–18
 - public vs. undocumented, 16–17
 - third-party, 365
- AppInfo!AilsEXESafeToAutoApprove function, 499
- AppInfo!RAiLaunchAdminProcess function, 499
- AppInfo service, 498
- application processes, 10–13
- applications. *See also* software
 - blocking from exiting on crash, 125
 - dispatcher objects, waiting on, 170–171
 - memory usage and, 473–474
 - program instructions, tracing, 153
- Application Verifier, 188, 206–217
 - downloading, 206–207
 - hooks, 210–214
 - hooks, enabling, 311
 - memory corruption bugs, catching, 331–332
 - operating system support, 209–214
 - page-heap configuration settings, 279–281
 - page heap, enabling, 272, 275–276
 - resource leaks, detecting, 307–310
 - systemwide settings, 217
 - verifier checks, enabling, 207
 - verifier check settings, 213
 - WinDbg and, 206
- arguments, saving in registers, 52
- ARM processors, 5
- artificial deadlocks, 239, 253
- assemblies, managed, 25
- assembling, 263
- assembly-mode stepping, 52
- _ASSERT macro, 138
- assertions
 - compile-time, 138
 - run-time, 138–139
 - system code, 6
- _ASSERT macro, 138
- assumptions, validating, 138
- asynchronous procedure calls (APCs), 15
- asynchronous processing, 342
- asynchronous UI events, 374–381
- AtIA2WHelper function, 300
- ATL (Active Template Library), 22, 29
- at (@) sign, 178
 - @! character sequence, 180–181
- attach operations, wait loops for, 220–222
- Attach To Process menu action, 160
- authentication, user, 10
- authorization rules, 354
- AutoLogger registry key, 449
- !avr command, 214–217, 276
 - symbols search path settings, 214–215

B

- background tasks, 10
- ba command, 162, 510, 525
 - /p option, 167
- bad format string bugs, 196–199
- bad resource handles, 206
- bang commands, 59. *See also* specific command names
- bc command, 165, 509, 526
- bd command, 80, 509–511, 526
- BeingDebugged field, overwriting, 259–260, 517
- bins, registry, 173
- Blaster virus, 201
- bl command, 509–510

- blocked time
 - analyzing, 458–473, 503
 - causality chains, reconstructing, 471–472
 - context switch events, 459–461
 - PerfView analysis, 482
 - wait analysis in PPA, 461–467
 - wait analysis in Xperf, 467–474
 - block header structure, 278
 - StackTrace field, 278
 - bm command, 510
 - boot tracing, 258, 449–454
 - kernel provider events, 450–452
 - Processes Summary Table, 451
 - Run and RunOnce applications, 451
 - user provider events, 452–454
 - bottlenecks, identifying, 402–405
 - bounds checking, 201
 - bp command, 40, 81–82, 224, 509–510, 525
 - /p option, 81
 - !bpmd command, 109, 111
 - breaking
 - on access, 162, 510
 - at custom action entry point, 243
 - on DLL load events, 222–227, 260–261
 - on first-chance exception notifications, 252–253
 - on nt!PspInsertProcess function calls, 255–257
 - on unhandled exceptions, 262
 - on user-mode process startup, 255–258
 - break-in requests, 221
 - break-in sequence
 - host debugger state, 102
 - kernel-mode, 77–78, 99
 - loader lock and, 352
 - user-mode, 91–93
 - break-in threads, 91
 - BreakOnDllLoad value, 225–228, 260–261
 - deleting, 261
 - breakpoint memory addresses, interpretation, 81
 - breakpoints
 - clearing, 526
 - code. *See* code breakpoints
 - conditional, 257–258
 - data. *See* data breakpoints
 - disabling, 526
 - at DllMain entry point code, 224
 - DLL mapping, setting at, 222
 - hardware, 163–164
 - in kernel-mode code, setting, 524
 - at nt!PspInsertProcess function, 255–257, 366
 - scope, restricting, 257–258, 382
 - source-level, 225
 - unresolved, 510
 - in user-mode code, setting, 525
 - bu command, 510
 - buffer overruns
 - catching, 199–201
 - exploitation of, 201
 - guard page position and, 280
 - heap-based, 275–277
 - immediate access violations, 281
 - stack-based, 291–293
 - buffers, inserting into memory locations, 97
 - buffer underruns, 280
 - build flavors, operating system, 5–6
 - built-in debugger commands, 59
 - busy-spin loops, 347
 - bytes, editing, 259
- ## C
- cabling protocols, 67–69
 - memory dump generation speeds, 156–157
 - caching file reads, 403
 - callbacks, canceling, 336
 - callers and functions, contracts between, 199–202
 - call frames
 - memory, displaying, 299
 - navigating between, 514–515
 - reconstructing, 295–297
 - symbolic names, 278
 - calling conventions for Win32 functions, 49–52
 - call stacks
 - allocating space, 300
 - collecting, 432–433
 - creation, 495
 - deep, 512
 - displaying, 511–512
 - dumping, 110
 - frame numbers, 48, 511
 - function parameters, displaying, 512
 - guard location, 291–292
 - handle-tracing, 305–307
 - incomplete, 439–441
 - kernel-mode side, 65–66
 - layout in __stdcall calling convention, 50
 - managed frames, 109
 - missing, 436–437
 - .NET call frames on, 440–441

call stacks

- call stacks, *continued*
 - parameters and locals, finding, 49–52
 - parameters, displaying, 48
 - of readying threads, 470–471
 - relative addresses in, 305
 - standard calling convention, 291
 - of thread waits, 465–467, 469
 - unloaded code and, 341–342
 - unresolved, 437–439
 - walking, 50–51
- CancelKeyPress event, 374–375
- capture anywhere/process anywhere paradigm, 399, 411
- C_ASSERT macro, 138
- causality chains, 457
 - reconstructing, 471–472
- c command-line option, 177, 230
- C/C++ applications
 - manual thread creation, 330–331
 - OACR static-analysis system, 202–204
 - verifier checks for, 207
- C/C++ development
 - debugging, 39–43
 - fields, displaying, 513–514
 - multiple overloads, 57–58
 - WDK environment for, 29
- C/C++ linker, 45
 - PDB files, 53
 - symbol decoration, 57
- __cdecl calling convention, 50
- cell indexes, 173
- cells, registry, 173
 - address, 174
- .chain command, 59
- checked builds, 5–6
- CheckToBreakOnFailure function, 249
- .childdbg command, 235, 238
- child debugging, 234–244
 - advantages and disadvantages, 239
 - enabling, 234–235, 243
 - functionality, 237–238
 - MSI custom actions, 240–244
 - startup code path debugging, 238–239
- child partitions, 74
- child processes
 - starting, 11
 - terminating, 11–13
- Chk* macros, 250–251
- circular logging, 421–422
- class factory objects, 21
- class identifiers (CLSIDs), 20, 24
- client operating systems, 4
- client/server runtime subsystem process. *See* csrss.exe
- CLR (Common Language Runtime), 25–28
 - current version, displaying, 108
 - debugging objects, 104
 - ETW events, 422
 - ETW instrumentation, 482
 - execution engine DLL, loading, 107–108
 - headers, 26
 - heap, 282
 - heap corruption and, 271
 - JIT compiler, 25, 103, 125
 - loading, 27
 - managed-code debugger contract, 103–104
 - managed debugging assistants, 289–291
 - mark-and-sweep scheme, 282
 - safety guarantees, 281
 - SOS debugger extension, 106–107
- !clrstack command, 109, 111, 284
 - a option, 110
- CLSIDs (class identifiers), 20, 24
- cmd.exe
 - leaked instances, 139
 - security context, 360
- code. *See also* C/C++ development; C++ development, debugging during, 39
 - instrumenting, 391, 411
 - instrumenting, fprintf statements, 416
- code analysis
 - runtime, 206–217
 - static, 195–206
- code base unification, 3
- code breakpoints
 - addresses relative to current process, 166
 - deactivating, 509–510
 - decorated names, setting by, 57
 - deleting, 43, 509
 - vs. execution data breakpoints, 166–167
 - hardware exceptions and, 89
 - hexadecimal addresses, setting by, 57
 - inserting in kernel-mode debugging, 81–83, 100
 - inserting in user-mode debugging, 93
 - insertion, observing, 93–97
 - listing, 509–510
 - live kernel debugging and, 67
 - at nt!PsInserProcess function call, 255–257
 - on paged-out memory, 100
 - at program entry point, 40–41

- relative to process context, 100
- setting, 40, 80, 81–83, 509–510, 524, 525
- symbolic name of command, matching, 510
- target location, 97
- unresolved, 510
- virtual addresses, 45
- code examples
 - dollar sign (\$) in, 40
 - ellipsis (...) in, 41
- COM (Component Object Model), 19–25
 - activation, 23–24
 - crosslanguage interoperability, 21
 - frameworks, 21
 - interprocess communication protocol, 19
 - language tools, 21
 - library, 21
 - module lock counts, 22
 - objects, consuming, 23–25
 - objects, writing, 21–23
 - out-of-process activation, 238–239
 - SCM, 21, 23
- COM Interop, 27
- command prompts, security context, 360–361
- commands. *See also* specific command names
 - history, recapturing, 178
- command window, docking, 41–42
- Common Language Runtime (CLR). *See* CLR (Common Language Runtime)
- communication mechanisms, 13–16
 - interprocess, 15–16
 - kernel-mode to user-mode, 15
 - user-mode to kernel-mode, 13–14
- companion source code, 29. *See also* scripts
 - boot tracing script, 450, 452
 - compiling programs, 39
 - error-handling macros, 248–250
 - ETW tracing scripts, 398
 - kill.exe emulation, 491–492
 - trace file viewing script, 399
 - trace symbols caching script, 408
 - XML file creation tool, 481
- compile-time assertions, 138
- complete memory dumps, 155
 - generating, 156
- COMPLUS_GCStress setting, 287
- COMPLUS_HeapVerify variable, 285, 287
- COMPLUS_Mda variable, 290
- Component Object Model (COM). *See* COM (Component Object Model)
- “Component Object Model Specification”
 - document, 19
- compressed data, 474
- COM serial cables, 68
- COM servers
 - in-process, 24
 - interface versioning, 26
 - lifetime, managing, 22–23
- Concurrency Analyzer, 461. *See also* PPA (Parallel Performance Analyzer)
- conditional breakpoints, 257–258
 - ignoring, 258
- consent.exe, 496–497
- console applications
 - asynchronous UI events, 374–381
 - Ctrl+C signal handling, 374–380
 - debugging, 370–373
 - drawing area management, 367, 373
 - functionality, 365
 - I/O functions, 368–373
 - open resources disposal, 377
 - printing text to screen, 366–373
 - termination handlers, 375–377
 - UI message-loop handling, 367
 - UI processing, 367
- console, displaying strings to, 365–380
- console host process (conhost.exe), 366–369
 - creator of, 367
 - I/O functions, 368
 - multiple instances, 370–371
 - UI events handling, 373–374
 - window drawing area management, 367, 373
- console subsystem, 365–380
- const modifier, 270
- constructors, calling, 169
- consumers, ETW, 417
- context record structures, 144
- context states, sharing with new threads, 330
- context switches, 92, 101–102, 522–523
 - analyzing, 468–469
 - conditions for, 459
 - CPU, 186
 - invasive, 81, 100, 101–102, 523
 - logging, 424
 - to SEH exception context record, 145
 - single-stepping and, 100–101
- control flow commands, 507–508
- controllers, ETW, 417
- cookie value, global, 292
- copy-on-write OS mechanism, 166

- C++
 - compile-time assertion mechanism, 138
 - expressions, evaluating, 278
 - global and static objects constructors and destructors, calling, 169
 - global objects initialization, 351
 - global state destruction, 337
 - reference-counting class, 335
 - template names, resolving, 180–181
- @@C++ prefix, 258
- CPU
 - analyzing usage, 404–408
 - architectures, 4–5
 - compressed data in, 473–474
 - instructions, use of, 8
 - scheduling, 459–461, 468–469
 - tracing usage, 397–398
 - viewing threads, 464
- CPU-bound operations, 397–405, 409, 502
- CPU context, switching, 186, 519
- CPU interrupts, 91
- CPU registers, 164
- CPU sample profiling, 404–406, 410, 458
 - bottlenecks, finding, 502
 - event logging, 424
 - PerfView for, 482
 - viewing, 464
- CPU Sampling By Process graph, 405
- CPU Sampling Summary Table, 406–407
- CPU Scheduling graph, 468–469
 - columns and descriptions, 469
- CPU trace flag, 101
- .crash command, 155
- crash dumps, 139. *See also* dump debugging; dump files
 - analyzing in Visual Studio debugger, 150–151
 - analyzing with WinDbg, 143–150
 - generation, automatic user-mode, 139–143
 - generation, manual, 151–153
 - generation settings, 154
 - kernel-mode, 154
 - loading, 57
 - LocalDumps registry key, 140
 - managed-code, 145–149
 - saving, 155–156
 - storage location, 140
 - types of, 140
- crashes
 - access violations, 267–270. *See also* access violations
 - automatic capture, 125
 - breaking in at, 127–128
 - exceptions leading to, 143–145
 - reproducing, 198
 - user-mode, 126, 262
- crashing processes
 - inspecting, 129–131
 - PIDs, 142
- CreateProcess API, 270
 - arguments, finding, 49–52
- CreateService.cmd script, 320
- CreateThread API dwStackSize parameter, 297
- creation call stacks, viewing, 495
- critical path of performance delays, 461
- critical sections
 - dumping, 345–347
 - leaks from, 206
 - usage bugs, 206
 - user-mode debugging support, 347
- C runtime library (CRT)
 - C++ program generic entry point, 168–169
 - I/O functions, 369–373
 - printf function implementation, 366
 - process rundown, 336–337
- !cs command, 345–347
- cscript.exe, 112
 - //X command-line option, 114, 135
- C# code
 - execution of, 385–386
 - unsafe, 281–283
- csrss.exe, 10
 - conhost.exe instance creation, 367
 - event handlers, injection, 374–380
 - kernel debugging, 377
 - UI message-loop handling, 367
- CSwitch events, 410, 459–461, 464, 491
 - filtering, 468
 - viewing, 498
- CSWITCH flag, 424, 459
 - adding to Base group, 467
 - enabling, 495
 - stack-walking flags, 432
- Ctrl+Alt+Break shortcut, 221
- Ctrl+Alt+Del shortcut, 9
- Ctrl+Alt+D shortcut, 76
- Ctrl+Break shortcut, 72, 375
- Ctrl+C shortcut
 - handling, 374–380
- Ctrl+Enter shortcut, 196
- current address, jumping to, 262–265

- current function, disassembling, 92
- current process
 - breakpoint scope, restricting to, 382
 - context switches, 81, 101–102
 - displaying, 520
- current process context alias, 179
- current thread context
 - after break-in, 91–92
 - checking, 102
 - switching, 46–47
- custom actions, debugging, 240–244
- custom allocators, 271
- cyclic dependencies, 351

D

- dangling threads, 337–338
- data breakpoints, 162–176
 - at global object destructor location, 169
 - clearing, 165–166
 - vs. code breakpoints, 166–167
 - debug registers, 164
 - disabling, 511
 - drawback of, 167
 - execution, 166–167
 - implementation, 163–165
 - in stack corruption investigations, 294–295
 - kernel-mode in action, 170–172
 - number of, 166
 - registry values, monitoring changes, 173–176
 - scope, restricting, 167
 - setting, 162–163, 510–511, 525–526
 - signal state of objects, watching, 171–172
 - survival of, 165
 - user-mode in action, 168–170
- db command, 97
- dbghelp!MiniDumpWriteDump API, 152
- DbgManagedDebugger registry value, 135
- dbgsvr.exe, 119
- .dbg symbol files, 53
- DCOM (Distributed COM), 20, 24
- DComLaunch service, 21, 238–239
- d* commands, 516–517
- dd command, 51, 59
- DDK (Windows Driver Development Kit), 28–29
 - build environment, 6, 203
- //D (double-slash) command-line option, 221
- deadlocks
 - analyzing, 156
 - artificial, 239, 253
 - !cs command, 345–347
 - debugging, 343–352
 - defined, 343
 - loader lock, 351–352
 - lock-ordering, 344–348
 - logical, 348–352
 - prevention, 347–348
 - threads, identifying, 345
- Debug\Break All menu action, 221
- Debug\Break menu action, 45, 72, 92
- DebugBreakProcess API, 91
- debug breaks
 - on HRESULT failures, 248
 - raising, 135–137
 - script-specific, 137
 - verify_heap method, 285
- debug builds, 5
- debug control register, 164
- _DEBUG_EVENT data structure, 237–238
- debug events, 88–91
 - filter states, resetting, 253
 - from multiple processes, 237–238
 - processing, 87–88
 - receiving and responding to, 86
 - types, 88
 - waiting for, 87, 88
- /debugExe option, 230
- debugger commands. *See also* specific command names
 - documentation of, 58–60, 83
 - kernel-mode, 63, 519
 - output, 41
 - scripting, 176–178
 - types, 59–60
 - user-mode, 519
- debugger comments, 40–41
- debugger extensions, 59
- debugger hooks, 187–194
 - BreakOnDllLoad IFEO hook, 260–261
 - DLL load event hook, 225–227
 - HRESULT failures hook, 248–250
 - Image File Execution Options hooks, 193–194
 - ntdll!g_dwLastErrorToBreakOn hook, 245
 - NT global flag, 187–193
 - startup debugger hook, 227–234, 239
- debugger loop
 - kernel-mode, 99
 - user-mode, 87–88
- debugger runtime controller thread, 103–104

debuggers

- debuggers. *See also* kernel-mode debuggers; user-mode debuggers; Visual Studio debugger; Windows debuggers
 - attaching to programs, 44
 - attaching to target, 220–222, 505
 - automated stepping, 251
 - execution environment, changing, 491–493
 - local and remote instances, quitting, 234
 - preventative use, 195
 - source file resolution, 41
 - thread suspension, 254–255
- debugger statement, 137
- debugging
 - access-check failures, 352–362
 - access violations, 267–270
 - automating, 178–183
 - COM failures, 21, 24
 - common tasks, 219
 - CPU architecture and, 5
 - custom action code, 240–244
 - deadlocks, 343–352
 - DLL load events, 222–227
 - error-code failures, 245–252
 - handle leaks, 300–307
 - heap corruptions, 271–291
 - high-level errors, 248–250
 - kernel-mode, 60–83. *See also* kernel-mode debugging (KD)
 - kernel-mode memory leaks, 316–321
 - live kernel, 61–67
 - live production environments, 33, 39
 - low-level errors, 247–248
 - modes, 52
 - noninvasive, 159–161
 - process startup, 227–234
 - race conditions, 323–343
 - source-level, 52–53. *See also* source-level debugging
 - debugging
 - stack corruptions, 291–297
 - stack overflows, 297–300
 - storage of information, 53
 - system code, 44–47
 - third-party applications, 33
 - tracing and, 490–502
 - unloaded code, 341–342
 - user-mode, 39–60. *See also* user-mode debugging
 - debugging
 - user-mode memory leaks, 307–316
 - Win32 debugging APIs, 87–88
 - Windows components, 33
 - WOW64, 183–187
 - your own code, 39–43
- debugging sessions
 - kernel-mode, ending, 526
 - kernel-mode, starting, 519
 - quitting, 225
 - terminating, 43
 - user-mode, ending, 518
 - user-mode, starting, 505
- Debug\Options And Settings dialog box, 132
- debug port executive object, 98
- debug port kernel object, 86
- debug port objects, 88
 - sharing, 237
- debug registers, 164
 - resetting, 166
- Debug\Source Mode WinDbg UI menu action, 53
- debug status register, 164
- defense in depth, 360, 495
- delegates, .NET, 376
- deleted memory, access attempts on, 268
- dependencies, cyclic, 351
- destructors, calling, 169–170
- developer interface, 16–28
 - CLR, 25–28
 - COM, 19–25
 - KMDF, 17
 - NTDLL module, 18
 - UMDF, 17–18
 - USER32, 18
 - WDM (Windows Driver Model), 17
 - Win32 API layer, 18–19
- developer tools, 28–29
- development kits, 28–29
- devenv.exe, 121. *See also* Visual Studio debugger
- device driver stacks, 14
- DeviceIOControl API, 319
- DisableAuto.bat command, 128
- DisablePagingExecutive registry value, 440–441, 464
- DisableThreadLibraryCalls API, 351
- disassemblies
 - displaying, 51, 212, 515–516
 - inspecting, 5
- disk access speed, 403
- disk I/O
 - analyzing, 403
 - event logging, 397
 - performance issues, 402
 - speed of, 473

- DISK_IO_* flags, 425
 - stack-walking events, 432
- Disk Utilization graph, 402–404
- dispatcher events, logging, 424
- DISPATCHER flag, 424, 460, 464
 - adding to Base group, 467
 - stack-walking events, 432
- dispatcher objects, 170–172
 - header structure, 171
- Distributed COM (DCOM), 20, 24
- DllCanUnloadNow functions, 22
- DLL debugger hooks, 193–194
- DLL hell, 26
- DLL hosting modules, 21
- DLL load events
 - breaking on, 222–227, 260–261
 - handling of, 252
- DllMain
 - entry point code, breaking at, 224
 - loader lock on, 351
 - synchronizing calls to, 351
- DLL modules
 - debugging in Visual Studio debugger, 225–227
 - entry point wrapper, 336
 - initialization work, 351
 - leaks, 307–308
 - lifetime-management bugs, 340–343
 - loaded, listing, 508–509, 523–524
 - loading and inspecting, 57
 - loading dynamically, 224
 - pinning, 343
 - reference counting, 342–343
 - unloaded, access violations, 268, 324, 340–343
- DLLs
 - binding and versioning, 26
 - custom-action, 240–244
- !dlls command, 524
- DLL_THREAD_ATTACH notifications, 351
- .dmp extension, 150
- Dock menu action, 41
- !do command, 110, 111, 149, 286
- dollar (\$) sign, 40, 178
- //D option, 137
- dot debugger commands, 59. *See also* specific
 - command names
- double-free bugs, 272
 - catching, 208
- dps command, 278–279
 - traversing stack with, 295–297
- driver development kits, 29
- driver extensions, 17
- drivers
 - interdriver communication, 14
 - kernel-mode, 17
 - registering, 320
 - signed, 8
 - writing, 17
- driver verifier tool (verifier.exe), 213
- !ldso command, 148, 149
- dt command, 48, 54, 171, 238, 513–514
 - r option, 80
- du command, 52, 246
- .dump command, 156, 160
 - /ma option, 152
- dump debugging, 139–157
 - analysis stage, 143
 - analysis with Visual Studio debugger, 150–151
 - analysis with WinDbg, 143–150
 - cause of crash, determining, 149–150
 - .dmp extension, 150
 - full memory dumps, 151
 - in-process managed debugging architecture
 - and, 105
 - kernel-mode, 153–157
 - manual crash-dump generation, 151–153
 - minidumps, 151–152
 - SOS extension for, 106
 - trace file analysis, 153
 - user-mode crash-dump generation, 139–143
- dumper action, 444
- dump files, 125, 139. *See also* crash dumps
 - bit masks, 152
 - comments, 152
 - full memory dumps, 151
 - manual generation, 151–153
 - minidumps, 151–152
 - name, date, type, and storage location, 140
 - overwriting, 152
 - saving, 152
 - snapshot information, 153
- !dumpheap command, 489
 - stat option, 488, 490
- dumping
 - access tokens, 360–361
 - block structure of memory allocations, 274
 - critical sections, 345–347
 - GC heap statistics, 488
 - handle leaks, 356
 - memory as sequence of function pointer
 - values, 278

dumping

- dumping, *continued*
 - objects, 110, 286
 - security contexts, 356
 - types, 238, 513–514
 - unicode strings, 246
- !dumpobj command, 111
- !dumpstack command, 110
- DuplicateHandle API, 305
- dv command, 48, 513
- DWORD values, editing, 245
- dwStackSize parameter, 297

E

- eax register, 252, 383, 385
 - function call return codes, 251
 - nt!PsInProcess parameters, 256
- eb command, 259, 265
- ebp register, 51
- e* commands, 517
- .ecx command, 144, 145
- ed command, 245
- edit_auto_logger.cmd script, 453
- edx register, 383
- !eeheap command, 488, 490
- !eeversion command, 108, 111
- .effmach command, 186
- elevated administrative command prompt, service termination and, 492
- ellipsis (...), 41
- EnableProperty registry value, 453
- entry point function, 336
- entry points
 - breaking at, 40–41
 - CRT-provided, 168–169
- EnumerateTraceGuidsEx API, 428
- _EPROCESS structure ImageFileName field
 - ImageFileName field, 256
- error-code debugging, 245–252
 - brute-force method, 251–252
 - ntdll!g_dwLastErrorToBreakOn hook, 245–247
 - Process Monitor for, 247–248
- error-code tracing, 490–494
- !error command, 210
- error-handling macros, 248–250
- esp register, 49, 384
- ETW (Event Tracing for Windows), 391, 415–456
 - administrative privileges, 419
 - architecture, 416–422
 - boot tracing, 449–454
 - call stacks, incomplete, 439–441
 - call stacks, missing, 436–437
 - call stacks, unresolved, 437–439
 - components, 417–418
 - consumers, 417
 - controllers, 417
 - existing Windows instrumentation, 422–431
 - flags and groups, 423–425
 - front-end UI tool, 391–413
 - improvements to, 411
 - in-memory buffers, 420
 - kernel events, stack tracing, 432–434
 - LOADER flag, 423
 - logging, enabling, 441–449
 - memory leaks, investigating, 310
 - NT Kernel Logger session, 418–419
 - PROC_THREAD flag, 423
 - providers, 417
 - sample profiling instrumentation, 445
 - sessions, 417, 419–422
 - stack-walk events, 431–441
 - user events, stack tracing, 434–436
 - Win32 APIs, 445–449
 - ETW event logging, 410, 445–449
 - circular, 421–422
 - CPU overhead, 416
 - in-memory buffers, 417
 - starting, 418, 482
 - viewing sessions, 421
 - ETW events, 441–445
 - CSwitch events, 410, 459–461
 - custom, 415
 - event descriptor header, 441–443
 - kernel, 496–499
 - marks, 430–431
 - Profile events, 410
 - ReadyThread events, 459–461
 - stack-walk events, 431–441
 - Start and End events, 445–449
 - types of, 401
 - user, 499–502
 - ETW heap trace provider, 476–480
 - ETW traces. *See also* traces; tracing
 - analyzing, 399–410
 - Base group of events, 397–398
 - collecting, 397–398
 - collecting with PerfView, 482–483
 - file location, 398
 - filtering events, 398

- graphs and data views, 402–405
 - for investigation overview, 490
 - kernel rundown information, 418, 438
 - lost events, 420
 - mark events, 430
 - merging with kernel rundown information, 483
 - Profile stack-walk switch, 397–398
 - scope of, 411
 - stack-walk events, 397, 405–407
 - starting and stopping, 398
 - system configuration information, 401
 - time interval, 400
 - viewing, 399
 - WPR/WPA use, 395
 - EventEnabled API, 448
 - event handlers
 - in .NET, 376
 - user-defined, 377, 379
 - event objects
 - signaled state, viewing, 80
 - viewing, 65, 79–80
 - EventRegister API, 448
 - events. *See also* ETW events
 - providers, 417
 - rate of incoming events, 420
 - schema for, 416
 - signaling, 170
 - EventUnregister API, 448
 - EventWrite API, 448
 - exception codes, 89
 - exception dispatching sequence
 - kernel-mode, 100
 - user-mode, 90
 - exception filter states, resetting, 253
 - exception handling, 88–91, 99
 - disabling, 223
 - exception records, 144
 - exceptions, 88–91
 - automatic crash-dump generation on, 140–143
 - breaking on, 262
 - inner and outer, 147–148
 - leading to crashes, finding, 143–145
 - nested, 146–147
 - .NET, 145–149
 - printing, 146
 - SEH, 143–145
 - stop on module load, 223–225
 - throw keyword, 89, 91
 - unhandled, 90
 - exception stops
 - enabling, 252
 - resetting, 253
 - executables
 - CLR headers, 26
 - main entry points, stopping at, 231
 - startup debugging, 227–234
 - execution data breakpoints, 162, 166–167
 - setting, 162
 - execution delays, 396–410
 - blocked time, 458–473
 - high-level cause, 397
 - hot spots, 397
 - investigating, 502–503
 - execution environment, altering with
 - debugger, 491–493
 - execution modes, 8–9
 - executive, 8
 - I/O manager component, 14
 - executive objects
 - address, finding, 64
 - creating, 255
 - executive service routines, entry points, 18
 - EXE hosting modules, 21
 - EXE-specific debugging hooks, 194
 - NT global flag, 187–193
 - existing debug port attach process, 161
 - extension (bang) commands, 59. *See also* specific
 - command names
 - extension DLLs, loading, 59
 - !extension_name.help command, 59
 - extensions
 - building, 16–28
 - driver, 17
- ## F
- F1 shortcut, 58, 83
 - F5 shortcut, 40, 226, 508
 - F6 shortcut, 44, 94, 120, 139, 160
 - F8 shortcut, 62
 - F9 shortcut, 225
 - F10 shortcut, 40, 42, 508
 - F11 shortcut, 42, 51, 508
 - __fastcall calling convention, 52
 - Fast System Call facility, 384
 - faults
 - register context, 144
 - SEH exceptions, 144
 - ~f command, 254

f command

- f command, 517
- File\Delete Workspaces menu action, 43
- FILE_IO_* flags, 425
 - stack-walking events, 432
- file object integrity levels, 361
- file reads
 - caching, 403
 - viewing, 402
- fill patterns of memory allocations, 273–274, 278
- FireWire cable, 69
- first-chance exception notifications, 90, 91
 - breaking on, 252–253
 - logging, 252
- flags
 - ETW, 423–425
 - stack-walking, 432
- .for command, 251
- FPO (Frame Pointer Omission), 439, 440
- fprintf statements, 416
- .frame command, 48, 514
- frame numbers, displaying, 48
- Frame Pointer Omission (FPO), 439, 440
- frame pointers, 50
 - displaying, 50–51
- framework-level exceptions, 88
- free builds, 5–6
- FreeLibraryAndExitThread API, 342
- free time, catching memory corruptions at, 273, 281
- freezing threads, 327
 - in kernel mode, 262–265
 - in user mode, 253–255
- fsutil.exe, 396
 - sparse files, 404
- full memory dumps, 151
- full page heap, 273
 - memory consumption, 280
- function addresses, resolving, 509
- function arguments, checking, 96
- function calls
 - disassembly of, 382
 - history, saving, 153
 - listing, 47
 - parameters and locals, finding, 49–52
 - return codes, 251
 - stepping over and into, 42, 251–252
- function disassemblies, viewing, 515–516
- function parameters, displaying, 512–513
- function pointers
 - memory, dumping as sequence of, 278
 - overwriting, 201

- termination handlers, 375–377
- functions
 - callers, contracts between, 199–202
 - no-inline, 249
 - referencing, 66
 - unassembling, 212, 382
- FxCop, 204–206
 - installation, 204–205

G

- g [address] command, 508
- garbage collection (GC), 25, 282
 - GC-heap validation at, 285–287
 - managed to unmanaged transitions, 27–28
 - thread suspension during, 282
 - triggering, 288–290
 - viewing events, 485
- gc command, 258
- GC Events By Time table, 485
- GC heap, 282
 - generations, partition by, 282
 - last “good” object, 286
 - live object roots, finding, 489
 - mark-and-sweep scheme, 282
 - memory usage, analyzing, 481–482
 - memory usage, analyzing with PerfView, 482–488, 503
 - memory usage, analyzing with SOS
 - extension, 488–490
 - objects, traversing, 284
 - statistics on, dumping, 488–489
 - validation at garbage collection, 285–287
- g command, 40, 45, 72–73, 95, 508
- !gcroot command, 489–490
- GCStress CLR configuration, 287–288
- gdi32.dll, 18
- GDR (General Distribution Release) patches, 7
- Generic Events graph, 426, 429–430, 436, 500
 - Start and End events, 448–449
 - Summary Table view, 430
- !gflag command, 191–192
- GFLAGS (gflags.exe), 188–191
 - native and WOW64 registry views, editing, 191
 - page heap, enabling, 275
 - startup debugger configuration, 227–229, 232
 - user-mode stack trace database, creating, 311–312
 - verifier bits, enabling, 216

- verifier hooks, enabling, 210–211
- g_hResultToBreakOn variable, 250
- !gle command, 59
- global cookie value, 292
- global flag bits, 192
- GlobalLogger registry key, 449
 - deleting, 452
 - editing, 450
- global objects, initialization, 351
- global stack trace database, 311, 314–316
- global state destruction, 337
- global variables
 - addresses, resolving, 509
 - modifiers on, 270
 - in worker threads, 336–339
- gn command, 376
- /GS compiler flag, 291–292
- guard pages, 273
- gu command, 508

H

- HAL (Hardware Abstraction Layer), 8
- !handle command, 59, 96–97
- handle leaks
 - debugging, 300–307
 - dumping, 356
 - example, 300–302
 - !htrace command, 302–307
 - verifier hooks, 308
- handles, 8
 - duplication of, 305
 - invalid, 209–211
 - properties, dumping, 304
- handle tracing, 302–306
 - disabling, 306
- hangs. *See* deadlocks
- Hardware Abstraction Layer (HAL), 8
- hardware-assisted virtualization technology, 74
- hardware breakpoints, 163–164
 - clearing, 165–166
 - setting, 510–511, 525–526
- hardware drivers, user-mode execution, 17–18
- hardware exceptions, 89
- HashAlgorithm.HashFinal method, 326
- HashAlgorithm.HashUpdate method, 326
- hash function chaining mechanism, 326
- hash objects
 - assigning to threads, 328
 - sharing, 326
- hash operations, locking, 328
- heap-based underruns and overruns, 272
- !heap command, 277–279
- heap corruptions, 206
 - debugging, 191, 193, 271–291
 - debugging complications, 271
 - !heap command, 277–279
 - managed, 281–291
 - native, 271–281
 - .NET code and, 271, 281
 - page heap, 272–277
- heap memory consumption, analyzing, 475–476
- Heap Outstanding Allocation Size graph, 479–480
- Heap Total Allocation Size graph, 478
- Heap Trace Provider, 476–480
- heap tracing, 475–480
 - allocation types, 480
 - analyzing with Xperf, 478–480
 - disabling, 478
 - enabling, 476–477
 - GC heap analysis, 481–482
- HeapVerify CLR hook, 285, 288
- .help command, 59
- Help files, debugger commands in, 58–60, 83
- hexadecimal addresses, setting breakpoints by, 57
- .hh command, 58, 83
- high-level errors, debugging, 248–250
- historical debugging, 153
- hives, registry, 173
- host debugger machines, 61, 67
 - administrative privileges, starting with, 70
 - communication with targets, 98
 - connecting to target, 70
- hosting modules, 21
- hot spots, investigating, 397–410
- HRESULT failure reporting, 248–250
- !htrace command, 302–307
 - diff option, 303, 305
 - disable option, 303
 - enable option, 302
 - functionality, 302–303
 - limitations, 307
 - snapshot option, 303
- Hyper-V, 73–74
- hypervisor boot driver, 73

I

- ia64, 4–5
- icacds.exe utility, 361
- IClassFactory interface, 19
- /i command-line option, 127
- IFE0 (Image File Execution Options) registry key, 189
 - BreakOnDllLoad value, 225–226
 - debugger hooks, 193–194
 - Debugger value, 228
 - TracingFlags value, 477
 - VerifierFlags value, 213
- IL, 25
- image load/unload events, logging, 424
- impersonation, 353, 354
- import table, updating, 212
- infinite loops
 - attaching debugger during, 220–222
 - suspending threads with, 262–265
 - unblocking, 221
- InitializeCriticalSectionAndSpinCount API, 347
- inline assertion routines, 138
- in-process COM servers, 24
- in-process debugging, 104. *See also* managed-code debugging
 - advantages and drawbacks, 106
 - script debugging, 113
- instruction pointer register (eip), 51
- instrumentation manifest
 - custom event payload, 449
 - location of, 427
- int 1 instruction, 93
 - single-stepping with, 100
- int 2c instruction, 138
- int 3 instruction, 91–92, 100, 138
 - breakpoints, setting, 93
- integrity levels, 358–360
 - UAC and, 495
- Intel, 4
 - Intel-VT processors, 73
 - Model-Specific Registers, 384
 - Xeon processors, 73
- IntelliTrace, 153
- interactive window station, integrity level, 359
- interfaces, documentation of, 16–17
- internal code invariants, confirming, 138
- Internet Explorer
 - JPEG GDI+ buffer overrun, 201
 - script debugging, enabling, 113
- interprocess communication, 19, 86, 98

- interthread dependencies, 458, 461
- invalid handle errors, 209–211
- invalid heap pointers, 272
- invalid memory access violations, 198, 267, 324
- invariants, broken, 200, 253, 323
- invasive context switches, 81, 100–102, 523
- I/O bottlenecks, 458
- I/O control commands (IOCTL), 13–14
- I/O manager
 - caching file reads, 403
 - performance optimizations, 403
- I/O processing, 380–386
 - in console applications, 369–373
 - viewing, 247–248
- I/O Request Packets (IRPs), 385
- I/O requests outstanding to disk, 402
- @\$ip pseudo-register, 180
- IRPs (I/O Request Packets), 385
- IUnknown interface, 19

J

- JIT (Just-in-Time) compiler, 25, 103, 125
- JIT (Just-in-Time) debugging, 125–139
 - attach operation, 125, 128–129
 - breaking in at crash site, 127–128
 - clearing debugger, 126
 - closing program without debugging, 129
 - CLR-specific JIT debugger, 135
 - functionality, 128–132
 - interacting with debugger, 139
 - invocation sequence, 131
 - managed-code, 133–135
 - run-time assertions and, 138–139
 - script debugging, 135–137
 - in session 0, 139
 - Visual Studio for, 132–137
 - WinDbg as default debugger, 127–128
- job objects, 11–13
- JPEG GDI+ buffer overrun, 201
- JScript, 112
- jscript.dll, 112
- jumps, short, 265
- jump to self instruction, 262–265
- Just-in-Time (JIT) compiler, 25, 103, 125
- Just-in-Time (JIT) debugging, 125–139

K

- KCB (Key Control Block), 173
- ~*k command, 92
- k command, 46, 47, 50, 59, 128, 511–512
 - frame count with, 512
 - stack chain corruptions and, 295–296
 - unloaded code and, 341
- k* commands, 160
- kd1394.dll, 77, 98
- kdcom.dll, 77, 98
- KD (kernel-mode debugging), 60–83. *See also* kernel-mode debugging (KD)
- kdsrv.exe, 119, 120
- KD transport extensions, 98–99
- kdbus.dll, 77, 98
- kernel, 8
 - client and server versions, 4
 - ETW providers, 422–426
 - handle tracing, 302–303
 - memory, 316–317
 - security of, 381
 - Windows NT, 3
- kernel32!CaptureStackBackTrace API, 314–315
- kernel32!CloseHandle API, 11
- kernel32!CreateProcessW API, 48
 - DEBUG_PROCESS flag, 238
- kernel32!CtrlRoutine function, 376–380
- kernel32!DeviceIoControl API, 14
- kernel32.dll module, 18
- kernel32!SetConsoleCtrlHandler API, 375–376
- kernel32!SetLastError API, 247
- kernel32!SetThreadContext API, 164
- kernel32!WaitForSingleObject API, 11
- kernel32!WriteConsoleA API, 299, 373
- kernel32!WriteProcessMemory API, 95
- kernel binaries
 - finding, 66
 - type information, 54
- kernel bug checks
 - analyzing, 153
 - crash-dump generation, 154
 - inducing, 155
 - rebooting on, 154
- kernel components, memory consumption of, 318
- kernel configuration manager (CM), 173
- Kernel Debugging dialog box, 70
- kernel debugging port, 69–70
- kernel_debugging_tutorial.doc, 83
- kernel dispatcher objects, 170–172
- kernel drivers, 16
- kernel flags, 422–425
 - stack-walk events associated with, 432
- kernel memory dumps, 154–155
- kernel mode, 8–9
 - layering structure, 8–9
 - system calls, 383–384, 385–386
 - user-mode code calling into, 380–386
 - user-mode code, invoking, 15
 - user mode, transitioning from, 10
- kernel-mode code
 - breakpoints, setting, 524
 - calling from user mode, 13–14
- kernel-mode debuggers
 - code breakpoints, setting, 81–83
 - driver-signing verification, disabling, 320
 - global scope, 81
 - unhandled exceptions, stopping on, 262
- kernel-mode debugging (KD), 60–83
 - architecture of, 98–99
 - break-in sequence, 77–78
 - cabling options, 67–69
 - capabilities, 60
 - code breakpoints, setting, 81–83, 100, 166
 - communication channel, 98
 - context switches of current process, 101–102
 - crash-dump analysis, 154–156
 - data breakpoints, 164, 165–166
 - debugger functionality, 98–102
 - debugger prefix, 73
 - ending sessions, 526
 - exception handling, 99
 - freezing threads, 262–265
 - high-level requirements, 98
 - host debugger machine, 61, 67
 - host/target communication issues, 76
 - live kernel debugging, 61–67
 - nt!PsplInsertProcess breakpoint, 366
 - over 1394 ports, 69–71
 - over the network, 67–68
 - packet-based transport protocol, 78
 - physical machines, setting up, 67–73
 - postmortem, 153–157
 - pseudo-registers, 178–179
 - quick start guide, 519–526
 - registry access events, monitoring, 172–176
 - remote stubs, 119
 - second machine, 61
 - setup, 61

kernel-mode debugging (KD)

- kernel-mode debugging (KD), *continued*
 - simulating control by user-mode debugger, 259–260
 - single-stepping the target, 100–101
 - starting sessions, 519
 - systemwide scope, 33
 - target, controlling, 78–80
 - target machine, 61–62, 67
 - tracing mode, 76
 - transport extensions, 98–99
 - tricks for, 255–265
 - use of, 33
 - user-mode process creation, 255–258
 - virtual machines, setting up, 73–76
 - Windows boot sequence, 258
 - WOW64 programs, 185–187
- kernel-mode driver framework (KMDF), 17
- kernel-mode drivers, 17
- kernel-mode memory
 - code breakpoints, setting, 82–83
 - debugging leaks, 316–321
 - pool tagging, investigating leaks with, 318–321
- kernel-mode stack, 297
- kernel modules, renaming, 66
- Kernel Patch Protection, 381
- kernel provider events
 - logging during boot, 450–452
 - stack traces for, 432–434
 - viewing, 425–426
- Key Control Block (KCB), 173
- kf command, 299–300
- kill.exe, 170
 - emulation of, 491–492
- KMDF (kernel-mode driver framework), 17
- kn command, 48, 514
- Knowledge Base articles, 17
- kP command, 48–49, 512–513

L

- language-level exceptions, 88
- large files, execution delays,, 396–410
- Large Text File Viewer, 444–445
- .lastevent command, 143, 223
- latency, 457
- latent bugs, 195
- LaunchProcess function, 270
- l command, 52–53
- leader smss.exe, 9, 10

- leaked cmd.exe instances, 139
- leaked resources, 309–310
- LeakRoutine routine, 319
- leaks. *See also* memory leaks
 - critical section, 206
- leased system threads, 102
- least privileges, 360, 495
- lifetime management, reference counting, 333–336
- light page heap, 273
- live kernel debugging, 61–67
 - administrative elevation, 63
 - command prompt, 63
 - debugger prefix, 73
 - limitations, 61, 67
 - managed process threads, observing, 104
 - memory dump generation, 151–153
 - startup code paths, 232–234
- live production environments, debugging, 33, 39
- lm command, 55, 66, 508
- ln command, 342
- .loadby command, 108
- .load command, 59
- loaded modules, listing, 508–509, 523–524
- LOADER flag, 423, 424
 - enabling, 495
 - stack-walking events, 432
- loader lock, 351–352
 - break-in sequence and, 352
- load events, stopping on, 223–225
- load exception handling, disabling, 223
- LoadLibrary function, 342
- local cache, symbols, 55–56
- local debuggers, 118
- LocalDumps registry key, 140
 - dump file type setting, 151
- Local Security Authority Subsystem (LSASS) process.
 - See also* lsass.exe
 - user providers of, 428–429
 - user rights computation, 353
- locals window, invoking
 - , 221
- LocalSystem account, 10
- local variables
 - displaying, 513
 - dumping, 48
 - listing values, 47–52
- lock contention, 317
 - execution delays, 458
- LockCount field, 347
- lock counts, 22

- lock-ordering deadlocks, 344–348
 - avoiding, 347–348
 - !cs command, 345–347
 - locks, 343. *See also* deadlocks
 - loader lock, 351–352
 - lock-ordering scheme, 347–348
 - orphaned, 206
 - log channels, OwningPublisher GUID and, 427
 - loggers, 417
 - logging
 - circular, 421–422
 - context switches, 424
 - ETW events, 410, 445–449
 - with fprintf statements, 416
 - kernel provider events during boot, 450–452
 - lock-free, 417
 - standardized framework for, 416
 - user provider events during boot, 452–454
 - verbose, 56
 - logical deadlocks, 348–352
 - logical waits, 410
 - logman.exe, 428, 443–444
 - logon/logoff
 - management, 9–10
 - user authentication, 10
 - LogonUser API call stack, 306
 - .logopen and .logclose commands, 178
 - low-level errors, debugging, 247–248
 - lsass.exe, 10. *See also* Local Security Authority
 - Subsystem (LSASS) process
 - access to, 358
 - user rights computation, 353
- ## M
- main entry points
 - breaking at, 260
 - stopping at, 231
 - main UI thread
 - stack trace of, 65
 - synchronous operations, scheduling on, 396–410
 - tracing performance, 490
 - malicious code, buffer overrun exploits, 201
 - managed code
 - JIT debugging, 133–135
 - loading, 26
 - native code, interoperability, 27–28
 - symbols path resolution, 123
 - managed-code debugging, 103–112
 - architecture for, 103–106
 - lack of native support, 106–107
 - sos.dll debugger extension module, 106–112
 - managed crash dumps, 150–151
 - managed debugging assistants (MDAs), 289–291
 - managed exceptions, printing, 328
 - managed-heap corruptions, 281–291
 - debugging assistants, 288–291
 - !VerifyHeap command, 282–288
 - managed objects
 - listing, 148–149
 - pinning in memory, 28
 - random access violations on, 282
 - managed-to-unmanaged code transitions, 288–290
 - mandatory integrity labels, 359, 361, 362
 - manifest-based user providers, 426–427
 - mark events, 430–431
 - marshaling, 183
 - ~m command, 255
 - MDAs (managed debugging assistants), 289–291
 - memory
 - buffers, inserting, 97
 - of call frames, 299
 - committed, 268
 - dumping as bytes, 51, 97
 - dumping as function pointer values, 278
 - freeing, 282
 - kernel, 316–317
 - locality effect, 474
 - mark-and-sweep scheme, 282
 - overwriting, 200. *See also* buffer overruns
 - virtual allocations, 268
 - memory access
 - speed of, 403, 473
 - violations, debugging, 267–270
 - memory addresses, jumping to, 262–265
 - memory allocations
 - aligned and unaligned, 273, 280
 - block structure, dumping, 274
 - call stack, viewing, 277, 278
 - fill patterns, 273–274, 278
 - global stack trace database, 311, 314–316
 - guard pages, 273
 - light- and full-page-heap schemes, 274
 - observing, 310
 - placement on heap, 280
 - tracing, 311–312
 - memory corruptions
 - free time, catching at, 273, 281
 - reproducibility, 199

memory dumps

- sign of, 331
- memory dumps, 139. *See also* dump debugging;
dump files
 - complete, 155, 156
 - kernel, 154–155
 - registry values, extracting, 172
 - small, 155, 156
 - user-mode, 151
- memory leaks, 307
 - allocation call sites, viewing, 479–480
 - Application Verifier, detecting with, 307–310
 - handle leaks, 300–307
 - kernel-mode, 316–321
 - outstanding allocations and, 479
 - stack-trace database, 314–316
 - tracking, 475–476
 - UMDH tool, investigating with, 310–314
 - user-mode, 307–316
- memory locations
 - alias to, 515–516
 - monitoring actions on, 162
- memory management, 25
 - overhead of, 457
- memory pages
 - access violations, 268
 - guard, 273
- memory pools
 - paged and nonpaged, 316–317
 - pool tagging, 318–321
- memory usage
 - analyzing, 457, 473–490, 503
 - debugger analysis vs. tracing, 489
 - in GC heap, 481–482
 - memory leaks, 476
 - of .NET programs, 474
 - in NT heap, 475–476
 - private bytes, 474
 - in target process, 474–475
- memory values
 - displaying, 516–518
 - editing, 517
 - filling range of, 517
 - overwriting, 517
- message notifications, 86, 90–91
- method invocations, 23–25
- MFC (Microsoft Foundation C++ Class library),
Microsoft Developer Network (MSDN), 16–17
- Microsoft Intermediate Language (MSIL). *See* MSIL
(Microsoft Intermediate Language) images
- Microsoft .NET Framework. *See* .NET Framework
- Microsoft private symbols, archiving, 54
- Microsoft public symbol files, 54, 406
- Microsoft public symbols server, 45, 54
 - adding to search path, 505–506
- Microsoft Visual Studio. *See* Visual Studio
- Microsoft-Windows-Services user provider
 - enabling, 452–453
 - graph of events, 454
- Microsoft Xbox 360, 4
- MIDL, 21
- minidumps, 151–152
- MIPS processors, 4
- mixed-mode debugging, 38
- Model-Specific Registers (MSR), 384
- ModLoad messages, 223
- module load events
 - handling of, 252
 - stopping on, 223–225
- modules
 - listing, 508, 523–524
 - loaded, listing, 55
 - symbols, listing, 45
- msconfig.exe
 - kernel-mode debugging, enabling, 61–62
 - safe mode, 62
- mscordbi.dll CLR debugging objects, 104
- mscoree!_CorExeMain method, 26
- mscoree.dll, 27
 - reverse COM Interop, 28
 - versioning, 26
- mscorlib.dll NGEN image, 109
- MSDN (Microsoft Developer Network), 16–17
- MS-DOS, 3
- MsiBreak environment variable, 244
- MSI custom actions, 240–244
- MSIL (Microsoft Intermediate Language) images, 25
 - debugging, 103–112
 - metadata, 107
 - PDB symbol files, 53
- MSR (Model-Specific Registers), 384
- msvcrt!exit function, 336–337
- msvsmon.exe, 121–123
- MultiByteToWideChar API, 300
- multiple overloads, 57–58
- multithreading
 - freezing threads, 253–255
 - race conditions, 323
 - visualizing, 462
- mutexes, 347

N

- named pipes, 73, 76
- native code
 - debugging, 38
 - garbage collections and, 289
 - managed code, interoperability, 27–28
 - .NET type consumption, 28
 - runtime analysis, 206
 - shared-state lifetime management bugs, 330–339
- native heap corruptions, 271–281
- native modules, 25
- ~n command, 255
- nested exceptions, 146–147
- .NET applications
 - debugger runtime controller thread, 103
 - debugging, 107–112
 - memory usage, 474
 - run-time code generation, 103
- .NET call frames, 440–441
- .NET call stacks, 458
- .NET classes, thread safety, 326
- .NET code
 - analyzing, 411
 - console termination handlers, 375–377
 - exception stops, enabling, 252
 - execution engine behavior, configuring, 289–291
 - FxCop static-analysis engine, 204–206
 - heap corruptions, 281
 - JIT debugging, 133–135
 - source-level debugging, 150–151
 - wait loops, attaching debugger during, 220
- .NET debugging, 38, 103–112
 - Visual Studio debugger for, 103
- .NET exceptions
 - finding, 145–149
 - retrieving, 148
- .NET Framework, 20
 - executable load order, 26–27
 - execution engine environment, 25
 - managed and native code interoperability, 27–28
 - side-by-side versioning, 26
 - SOS extension, 103, 106–112
- .NET objects
 - dumping, 148
 - layout, 286–287
- .NET SDK, 28
- network ports, displaying, 119
- NGEN'ing applications, 106
- _NO_DEBUG_HEAP environment variable, 193
- NoDebugInherit bit, 238
- no-inline functions, 249
- noninteractive sessions, 10
- noninvasive debugging, 159–161
 - detaching from target, 161
 - starting session, 160
- nonpaged memory pools, 316–317
- NoRefCountingBug.exe, 331
- notifications, 86, 90–91
- n!PspAllocateProcess function, 255
- NT_ASSERT macro, 138, 139
- ntdll!Csr* routines, 18
- ntdll!DbgUiRemoteBreakin function, 92
- ntdll.dll, 8, 18
 - NT executive services for transition to kernel, 380
 - thread pool implementation, 343
- ntdll!g_dwLastErrorToBreakOn variable, 245, 491
- ntdll!LdrpLoaderLock, 351
- ntdll!Ldr* routines, 18
- ntdll!NtDeviceIoControlFile, 14
- ntdll!NtTerminateProcess function call, 212
- ntdll!Rtl* routines, 18
- ntdll!RtlSetLastWin32Error function, 246
- nt!_EPROCESS, 255, 521
- nt!FunctionName notation, 66
- NT global flag, 187
 - bits in, 188, 191
 - editing, 188–191
 - nt!NtGlobalFlag variable, 189
 - per-process, 189–190
 - per-process, reading, 191
 - scopes, 187
 - stop on exception (soe) bit, 262
 - systemwide DWORD value, 189
 - systemwide vs. process-specific, 187–188
 - user-mode debuggers and, 193
 - verifier hooks, enabling, 211
- NT heap, 282
 - allocation call sites, 478
 - memory consumption, analyzing, 475–480
 - memory leaks, investigating, 310–314
- NT heap manager
 - allocators, 272
 - APIs in ntdll.dll, 271
 - buffer alignment, 273
 - memory allocation, 268
- NT heap provider
 - ETW instrumentation, 475
 - logged events, visualizing, 478–480

nt!KdCheckForDebugBreak frame

- nt!KdCheckForDebugBreak frame, 77
- NT Kernel Logger session, 418–419
 - buffer size, 420–421
 - enabling, 452
 - kernel rundown information, merging events with, 451, 453
 - starting and stopping, 419, 449
- nt!MmAccessFault, 100
- nt module, 66
- nt!NtGlobalFlag variable, 187, 189
 - current value, displaying, 192
- nt!poolhittag global variable, 321
- nt!PsplInsertProcess breakpoint, 255–259, 366
 - inserting, 260
- nt!SeAccessCheck function, 354
- nt!SeAuditProcessCreation function call, 256
- NT security model, 353–358
 - concepts of, 353–354
 - integrity levels, 358–360
 - Windows Vista improvements, 358–361
- NT service control manager, 10. *See also* SCM (service control manager)
- _NT_SYMBOL_PATH environment variable, 313–314, 408, 438
- _NT_SYMCACHE_PATH environment variable, 408, 438
- NULL pointer address, dereferencing, 268

O

- oacr.bat batch file
 - clean option, 204
 - set option, 204
- OACR (Office Auto Code Review), 202–204
- !object command, 65, 357
- object handles, 8. *See also* handles
 - closing, 300
- object pooling, 328–330
- objects
 - dumping, 110, 286
 - lifetime management, 333–336
 - pinned, 289
- o command-line option, 235, 238
- off-by-one bugs, 273
- Office Auto Code Review (OACR), 202–204
- ole32!CoFreeUnusedLibraries function, 22
- OLE Automation, 20
- orphaned locks, 206
- orphaned service processes, 234

- OS-level exceptions, 88
- OSR website, 18
- __out_ecount SAL annotation, 200
- out-of-process COM servers, 24–25
- out-of-process debugging, 104
 - advantages and drawbacks, 106
 - SOS extension support, 107
 - Visual Studio support, 106
- overloads, multiple, 57–58
- OwningPublisher GUID, 427

P

- page access violations, 268
- paged memory pools, 316–317
- page fault handler, 100
- page file, reserved storage in, 268
- page heap, 272–277
 - block header structure, 278
 - customizing, 279–281
 - DLLs, restricting scope to, 280
 - enabling, 272, 275–277
 - full, 273–274
 - functionality, 273–274
 - light, 273–274
 - light- and full-page-heap allocations, alternating between, 280
 - memory consumption settings, 280
 - memory impact, minimizing, 273
 - scope, restricting, 273–274
 - target process, enabling for, 275–277
- page heap bit, 191
- parallel execution, race conditions, 323
- Parallel Performance Analyzer (PPA). *See* PPA (Parallel Performance Analyzer)
- parameters
 - listing in your own code, 48–49
 - listing values, 47–51
- parent/child relationship
 - breaking, 234, 239
 - preserving, 240
- parent processes, 130
- PatchGuard, 381
- pc command, 508
- p command, 508
- .pdb (program database) file extension, 53
- PDB symbol files, 53–54
- pdm.dll, 112

- PDM (Process Debug Manager), 112
 - debugging support, enabling, 135–136
 - enabling, 221
- !peb command, 160–161, 236, 507
- PEB (process environment block)
 - BeingDebugged byte, 259–260
 - dumping contents, 160–161
 - process command line, 507
- @\$peb pseudo-register, 179
- \$peb pseudo-register, 178, 510, 517
- !pe command, 146–147, 149, 328
- pe command-line option, 161
- peer code reviews, 202
- performance
 - blocked-time analysis, 458–473
 - hot spot analysis, 405–407
 - investigating with Xperf, 502
 - issues, investigating, 397
 - memory usage and, 457, 473, 503
 - printf tracing and, 416
 - read-ahead optimization, 403
 - write-behind optimization, 403
- performance monitor (perfmon.exe) memory-usage inspection, 474
- PerfView, 481–482
 - allocations, viewing, 486–487
 - CPU sample profiling, 502
 - GC Events By Time table, 485
 - GC Heap Alloc Stacks view, 485, 488
 - GC heap memory usage, analyzing, 482–488, 503
 - GCStats view, 484–485, 488
 - heap snapshots, 488
 - [Just my app] filter, 486
 - process summary table, 483–484
 - starting, 482
 - symbols resolution, 486
- Perl, 112
- PHRACK magazine, 201
- P/Invoke, 27
 - incorrect declarations, 281
- pool tagging, 318–321
- !poolused command, 320
- postmortem debugging, 125–158
 - dump debugging, 139–157
 - JIT debugging, 125–139
 - kernel-mode, 153–157
- Power PC processors, 4
- PPA (Parallel Performance Analyzer), 458
 - blocked-time analysis, 503
 - Cores view, 464
 - CPU Utilization view, 464
 - ETW trace views, 464–465
 - Microsoft public symbols, access to, 461–462
 - Threads view, 464–466
 - tracing sessions, starting, 461–464
 - user interface, 458
 - uses of, 458
 - wait analysis, 461–467
- PRCB (processor control block), data breakpoints in, 164
- prerelease milestones, 6
- printf function, 366–373, 416
 - console host process (conhost.exe), 366–369
 - CRT I/O functions, 369–373
 - stepping through, 371–372
- printing exceptions, 146
- printing text to screen, 366–373
- private code, breakpoints in, 166–167
- private symbols, archiving, 54
- privileges, 353
- Process class, 493
- !process command, 64, 102, 129–130, 520–521
 - arguments of, 520–526
- .process command, 522–523
 - /i option, 81, 101–102, 167, 523–525
 - /p option, 101, 522
 - /r option, 65, 101, 522
- process creation
 - breaking at, 366
 - call stacks, 495
 - user-mode, debugging, 255–258
- Process Debug Manager (PDM), 112, 135
 - enabling, 221
- process environment block (PEB), dumping contents, 160–161
- processes, 10. *See also* target process
 - access tokens, 353
 - attaching to, 120
 - child processes, debugging, 234–244
 - code paths, tracing, 495
 - command line, 237
 - context switches, 522–523
 - debug port object sharing, 237
 - displaying information about, 520–521
 - execution, stepping through, 232–234
 - interprocess communication, 15–16
 - in job objects, 11
 - killing, 492–493
 - listing, 520

process

- process, *continued*
 - managing as groups, 11
 - parent/child relationship, 234, 239, 240
 - in session 0, debugging, 231–234
 - signaling, 170–172
 - start/exit events, logging, 424
 - terminating, 11–13
 - verifier checks, 217
- process global flag, 193
- process ID (PID)
 - checking, 97
 - of crashing processes, 130–131
- Process Lifetimes graph, 399–400, 423, 451
- Process Monitor, 172
 - access failures, observing, 247–248
- processor context switches, 519
- processor control block (PRCB), data breakpoints in, 164
- process rundown sequence, 336–339
 - dangling threads, 338
- process startup debugging, 227–234
 - in session 0, debugging, 231–234
 - user-mode processes, 259–260
 - during Windows boot sequence, 258
- Process Summary Table, 400–401
- @\$proc pseudo-register, 179
- \$proc pseudo-register, 377–378
- PROC_THREAD flag, 423, 424, 495
 - enabling, 495
 - stack-walking events, 432
- production environments, debugging, 39
- Profile events, 491
 - viewing, 498
- PROFILE flag, 424, 464
 - enabling, 495
 - stack-walking events, 432
- Profile stack-walk switch, 397–398
- programs, attaching debuggers to, 44
- protected page access violations, 268–271
- providers
 - ETW, 417
 - installation, 427
- pseudo-registers, 178–180
 - displaying value, 178
 - \$peb, 510, 517
 - strongly typed, 258
- @\$ptrsize pseudo-register, 180
- public symbols, 54
- pv command-line option, 160–161
- Python, 112

Q

- q command, 43, 518
- qd command, 43, 161, 518
- qq command, 118, 234
- quality assurance, 217
- quick start guides
 - kernel-mode debugging, 519–526
 - user-mode debugging, 505–518
- quitting and detaching from target, 518
- quitting and detaching without terminating target, 225
- quitting and terminating target, 225

R

- race conditions
 - deadlocks, 343
 - debugging, 323–343
 - defined, 323
 - DLL module lifetime-management bugs, 340–343
 - freezing threads, 253–255
 - object pooling, 328–330
 - reproducing, 323, 326–328
 - shared-state consistency bugs, 324–330
 - shared-state lifetime-management bugs, 330–339
 - simulating in debugger, 327–328
 - WaitForMultipleObjects API, 333
- RaiseException API, 89, 91
- @\$ra pseudo-register, 179
- \$ra pseudo-register, 178
- rax register, 251
- r command, 59, 166, 516, 518
- r? command, 258
- rdmsr command, 384
- read-ahead optimization, 403
- read data breakpoints, 162–163
- read-only memory, string literals in, 269
- readying threads, 460
 - displaying information on, 470
- ReadyThread events, 459–461
- ReadyTime interval, 460
- .reboot command, 258
- recovery, 170
- reentrancy issues, 416
- reference-counted objects, tracking, 314–316
- reference counting, 21–22, 333–336
 - for DLL module lifetime management, 342–343

- regasm.exe, 28
- !reg command, 172–176
 - openkeys option, 173
 - valuelist option, 174
- registers. *See also* pseudo-registers
 - arguments saved in, 52
 - Model-Specific Registers, 384
- register values
 - displaying, 516–518
 - overwriting, 518
- registry
 - changes, logging, 425
 - editing, 191
 - storage concepts, 173
- REGISTRY flag, 425
 - stack-walking events, 432
- registry values
 - changes, tracking, 172–176
 - editing, 188
 - Key Control Block, 173
 - listing, 173
- regression bugs, 217
- regular expression evaluation, 258
- releases, operating system, 6
- .reload command, 45, 55, 505–506
 - /unl switch, 341
 - /user switch, 524
- remote command-line option, 118, 232
- remote debugging, 116–123
 - architecture for, 116–117
 - starting, 118
 - in Visual Studio, 121–123
 - with WinDbg, 117–121
- Remote Desktop connection, kernel debugging
 - with, 75
- remote procedure calls (RPCs), 500
- remote sessions, 116
 - starting, 118
 - terminating, 118
 - in WinDbg, 117–118
- remote stubs, 116–117, 119–121
- resource handles, bad
 - , 206
- resource leaks, detecting, 307–310
- Resource Monitor (resmon.exe), 397
 - memory-usage inspection, 474
- .restart command, 59
- ResumeThread function, 254, 262
- retail builds, 5
- @\$retreg pseudo-register, 180

- reverse hazard, 337
- rights, user, 353
- RPCs (Remote Procedure Calls), client side, 500–502
- RpcSs service, 21, 24
- rsp register, 49
- _RTL_CRITICAL_SECTION structure, 346
- RtlEnterCriticalSection function, 345–346
- run-time assertions
 - inline assertion routines, 138
 - JIT debugging and, 138–139
- runtime code analysis, 206–217
 - !avr extension command, 214–217
- run-time code generation, 103
- run-time events, saving for analysis, 153

S

- safe mode, 62
- SAL (Standard Annotation Language), 195
- SAL annotations, 199–202
 - buffer overruns, 199–201
 - syntax, 199
 - for Win32 APIs, 202
- sample profiling, 404–406, 410
- scheduling, 10
- SCM (service control manager), 10
 - handshake timeout, 232–234
 - service entries list, 182
- s command, 92
- screen, printing text to, 366–373
- script debugging, 112–115
 - architecture for, 112–114
 - PDM, enabling, 221
 - startup code path, 220–221
 - in Visual Studio, 114–115
 - //X option, 114
- scripting, 176–183
 - command history, saving, 178
 - C++ template function names, resolving, 180–181
 - pseudo-registers, 178–180
 - replaying debugger commands, 176–178
 - startup debugging and, 230–231
 - Windows service processes, listing, 181–183
- scripting hosts, 112
 - enabling debugging, 114
- script JIT debugging, 135–137
 - debug breaks, 135–136
- script languages, 112

- scripts
 - attaching debugger, 114–115
 - blocked-time script, 458
 - boot tracing, 450, 452
 - ETW trace sessions, automating, 398
 - heap tracing configuration, 477
 - heap tracing session startup, 477–478
 - start_kernel_trace.cmd, 419, 423, 438
 - start_user_trace.cmd, 438
 - trace files, viewing, 399
 - trace symbols, caching, 408
 - view_trace.cmd, 438, 478
 - Xperf wrapper, 433
- !sd command, 357–358
- SDK ISO files, mounting and unmounting, 35, 37
- second-chance exception notifications, 90, 127
- security, 353–358
 - access tokens, displaying, 355–357
 - concepts of, 353–354
 - integrity levels, 358–360
 - UAC, 495
 - Windows Vista improvements, 358–361
- security checks, 352. *See also* access checks
- security context, 353
 - of command-prompt window, 360
 - dumping, 356
 - of user-mode debugger, 220
- security descriptors, 353, 354
 - displaying, 357–358, 360
 - reading and modifying, 361
- security identifiers (SIDs). *See* SIDs (security identifiers)
- SeDebugPrivilege, 87, 362, 493
- SEH (Structured Exception Handling) exceptions, 88–91
 - from assertion macros, 138
 - breaking on, 262
 - dereferencing NULL pointer, 126
 - exception codes, 89
 - finding, 143–145
 - handling, 89–91
 - hardware exceptions, 89
 - JIT debugging sequence, invoking, 138
 - memory access violations, 268
 - software exceptions, 89
 - Visual C++ language support, 89
- semaphores, signaling, 170
- serial COM ports, emulating, 73
- .server command, 117, 118
- server command-line option, 232
- server operating systems, 4
- service control manager (SCM). *See* SCM (service control manager)
- service packs, 7
- service routines, sharing between user and kernel sides, 381–382
- services
 - startup debugging, 231–234
 - status, reporting, 233
- services.exe, 10, 182
- services!ImageDatabase global variable, 182
- ServicesPipeTimeout registry value, 233
- service stubs, 18
- servicing updates, 7
- session 0
 - JIT debugging in, 139
 - startup debugging in, 231–234
- session manager subsystem process, 9
- sessions. *See also* debugging sessions
 - ETW, 417, 419–422
 - global and auto-logging sessions, 449
- session space, 9
- SetThreadPoolCallbackLibrary API, 343
- setup.bat file, 141–142, 232
- SHA-1 hash function, 326
- shared state
 - consistency bugs, 324–330
 - lifetime management bugs, 330–339
 - synchronization, 324, 326
- shell32!AicLaunchAdminProcess function, 501
- shell32!ShellExecuteW function, 501
- shellcode, 201
- Shift+F11 shortcut, 508
- short jumps, 265
- SIDs (security identifiers), 353
 - displaying, 355–356
 - friendly names, 357–358
 - integrity level, 358–359
 - owner, 362
- SignalState field, 171
- signal state of objects, 170–172
- signed drivers, 8
- single-stepping, 100–101
 - hardware exceptions and, 89
 - TF flag, 93
- small memory dumps, 155
 - generating, 156
- smart hosts, 112–113
- smss.exe, 9

- snapshots, 76
 - noninvasive debugging, 159
- soft-hang investigations, 156–157, 396–410
- software. *See also* applications
 - scalability and performance, 365
 - tracing and debugging, 365
- software breakpoints, 166. *See also* code
 - breakpoints
 - setting, 509–510
- software crashes, 125. *See also* crashes
- software development, 3–30
 - debugging during, 195
 - quality assurance, 217
- software exceptions, 89
- SOS extension, 103, 106–112
 - CLR architecture, matching to process, 146
 - debugger commands, list of, 111
 - GC heap memory usage, analyzing, 488–490
 - loading, 108
 - managed-code crash dump analysis, 146–149
 - !VerifyHeap command, 284–287
 - version, matching to CLR version, 146
- source code, navigating, 515
- source files
 - finding, 506–507
 - resolution of, 41
- source-level breakpoints, 225
- source-level debugging, 38, 52–53, 103
 - script debugging, 112–115
 - in scripts, 115
 - source path, 506–507
 - symbol files for managed assemblies, 107
- source-level .NET debugging, 150–151
- source-mode stepping, 52
- sources search path, 506–507
- sparse files, 404
- \$spat operator, 258
- .srcpath command, 41
- .srcpath+ command, 41
- SRM (security reference monitor), access check
 - initiation, 354
- SSDT (System Service Dispatch Table), 381
- STACK_COMMAND, 144
- stack corruptions
 - buffer overruns, 291–293
 - call frames, reconstructing, 295–297
 - data breakpoints, investigating with, 294–295
 - debugging, 291–297
 - k command and, 295–296
- Stack Counts By Type graph, 433, 494, 498
- Stack Counts Summary Table, 454
- stack-guard value
 - consistency check, 291–293
 - corruptions, investigating, 294–295
- stack objects, dumping, 148
- stack overflows
 - causes, 297–298
 - debugging, 297–300
 - kf command, 299–300
- stack pointer
 - restoring, 49–50, 383–384
 - saving, 383–384
- stack-trace database, 311, 314–316
 - snapshots of, 312–313
- StackTrace field, 278
- stack traces, 407, 409, 411
 - of access failures, 247
 - collecting, 443
 - current thread, 511
 - dumping, 65
 - of faulting instructions, 144
 - kernel-mode frames, 156
 - kernel provider events, 432–434
 - listing, 46
 - user provider events, 434–436
- stack-walk events, 405–407, 431–441
 - capturing, 397
 - enabling, 433
 - kernel provider, 432–434
 - payloads, 432
 - provider flags, 433
- stack-walking flags, 432
- Standard Annotation Language (SAL), 195
- start_kernel_trace.cmd script, 419, 423
- Startup And Recovery dialog box, 154
- startup debugger
 - advantages and disadvantages, 239
 - child/parent relationship, breaking, 234
 - configuring, 227–229, 232
 - functionality, 230–231
 - MSI custom actions and, 242
 - registry value, 194
 - resetting, 229
 - in session 0, 232–234
 - Visual Studio as, 230
- startup debugging, 227–234
 - child debugging and, 238–239
 - DLL startup code paths, 224–225
 - scripting debuggers and, 230–231
 - in session 0, 231–234

startup debugging

- startup debugging, *continued*
 - with startup debugger, 227–229
 - timeout issues, 232–234
 - user-mode processes, 255–260
 - Windows services, 231–234
- static_assert keyword, 138
- static code analysis, 195–206
 - running automatically, 196
 - SAL annotations, 199–202
 - standalone tools for, 202–206
 - Visual C++ for, 196–199
 - wscspy function, 275
- __stdcall calling convention, 49–52, 96
- stepping. *See also* single-stepping
 - assembly-mode, 52
 - through custom-action DLL code, 244
 - debugger commands, 508
 - into function calls, 251–252
 - over and stepping into, 42
 - through process execution, 232–234
 - through startup code path, 239
- STL (Standard Template Library), 29
 - source code, 506
- stop on exception (soe) bit, 192, 262
- stop on module load exceptions, 223–225
- stress testing, 326
- strings, displaying to console, 365–380
- Structured Exception Handling (SEH) exceptions.
 - See* SEH (Structured Exception Handling) exceptions
- suspend count, 254–255
- suspending threads, 254–255
 - infinite loops for, 262–265
- SuspendThread function, 254, 262
- svchost.exe instances, 181–183
- sx command, 252
- sxd command, 253
- sxd ld command, 223
- sxe command, 252
- sxe ld command, 108, 223–225, 243, 258
 - wildcard characters in arguments, 223
- sxr command, 253
- symbol decoration, 57–58
- symbol files, 53–54
 - caching, 408
 - loading, 55
 - for managed assemblies, 107
 - for MSIL binaries, 53
 - pointing to, 45
 - for traces, viewing, 399
- symbols
 - cache paths, 55
 - caching locally, 55–56
 - @! character sequence, 180–181
 - downloading, 406, 408
 - forcing reload of, 55
 - listing, 260, 261
 - looking up, 54, 505–506
 - managed-code path resolution, 123
 - mismatches, 56
 - missing or unresolved, 108
 - for OS binaries, 49
 - pre-downloading, 56
 - reloading, 65, 101, 505–506, 523–524
 - resolution issues, 56
 - resolving in PerfView, 486
 - UMDH tool, resolving in, 313–314
 - unloaded, reloading, 341
 - Xperf, loading in, 405–406
- symbols search path
 - fixing, 143, 505–506
 - fixing, script for, 176–177
 - setting, 45
 - verifier flags and, 214–215
- symbols server, 45, 53–54
- symcache action, 444
- symchk.exe utility, 56
- .symfix command, 45, 55, 313–314, 505–506
- !sym noisy command, 56
- .symopt command, 57
- .sympath command, 55, 314
- .sympath+ command, 506
- !sym quiet command, 56
- sync blocks, 286
- synchronization of shared memory, 324, 326
- synchronous operations, scheduling, 396–410
- syscall and sysret instructions, 380
- SYSCALL events, 491
 - tracing, 493–494
- SyscallExit events, 433–434
- SYSCALL flags, 425–426
 - stack-walking events, 432
- sysenter and sysexit instructions, 380, 383–384
- SYSENTER_CS_MSR (0x174) register, 384
- SYSENTER_EIP_MSR (0x176) register, 384
- SYSENTER_ESP_MSR (0x175) register, 384
- SysInternals tools, 17, 155, 247
- system APIs, catching misuses, 206
- system breakpoints, setting, 18, 44
- system calls, 380–386

- drawback of, 13
 - failures, tracing, 493–494
 - functionality of, 8
 - kernel-mode side, 385–386
 - locals and arguments, listing, 49–52
 - logging, 425
 - marshaling, 183
 - transition to kernel mode, 383–384
 - user-mode side, 381–383
 - user-mode to kernel-mode transitions, 13–14
 - system code
 - assertions, 6
 - debugging, 44–47
 - parameters and locals, listing, 49–52
 - system components, observing interactions, 60
 - System.Console class, 375–376
 - system DLLs, breaking on loading, 225–227
 - system internals
 - providers, progressive enablement of, 495
 - tracing, 494–502
 - system-level tracing, 503
 - system loader lock, 351–352
 - system memory, access to, 8
 - system process
 - Cid, 520
 - PID, 102
 - system processes, 9–10
 - user-mode debuggers, attaching, 239
 - system service calls, 8
 - System Service Dispatch Table (SSDT), 381
 - system services, 8
 - system threads, leased, 102
- T**
- target machines, 61, 67
 - communication with host, 98
 - connecting to host debugger, 70
 - kernel bug check, inducing, 155
 - kernel-mode debugging, enabling, 61–62
 - reboot after cabling setup, 70, 76, 77
 - remote debugging, 116–123
 - unblocking, 72
 - virtual machines, 73–76
 - target process
 - attaching debugger to, 87
 - blocking execution, 139
 - breakpoints, setting in kernel debugger, 81
 - child process creation, breaking on, 234–244
 - child processes, 88
 - continuing execution, 87, 88
 - control flow commands, 507–508
 - controlling, 86, 160
 - debugger, attaching to, 505
 - debug port object, 88
 - exceptions, 88–91
 - execution environment, 491–493
 - existing debug port attach, 161
 - freezing execution, 91
 - heap allocations, 478
 - heap tracing, enabling, 477
 - inspecting, 87
 - live memory dumps, 151–153
 - memory usage, analyzing, 474–475
 - page heap, enabling, 275–277
 - shutdown, 170
 - single-stepping, 100–101
 - starting, 87
 - starting under user-mode debugger, 230
 - startup command line, finding, 507
 - stop debugging, 87
 - terminating upon quitting, 225
 - threads, viewing, 92
 - user-mode debugger, attaching, 220–222
 - user-mode handle, 96
 - virtual address space, inspecting and editing, 86, 87, 95
 - task manager handle-count and kernel-memory-usage columns, 301–302
 - TCB (trusted computing base), 10
 - t command, 51, 251, 508
 - TCP/IP communication ports
 - opening, 118–119
 - releasing, 120
 - @\$teb pseudo-register, 179
 - \$teb pseudo-register, 178
 - TEB (thread environment block)
 - uninitialized fields, 352
 - Win32 error codes, 247
 - template functions, breakpoints in, 180–181
 - TerminateProcess API, 170, 212
 - termination
 - debugging sessions, 43
 - orderly or abrupt, 170, 172
 - termination handlers, console, 375–377
 - test-driven development, 39
 - testing
 - rare code paths and, 199
 - stress testing, 326

TF (trap flag) flag

- on virtual machines, 75–76
- TF (trap flag) flag, 93
- third-party applications, debugging, 33
- thrashing, 273
- !thread command, 65, 102, 521–522
- .thread command, 522
 - /r option, 186
 - /w option, 185–186
- thread context. *See also* context switches
 - 64-bit, switching to, 186
 - access token of, 355
 - debug registers in, 164
 - switching, 46–47, 522–523
- thread pools, 343
- @\$thread pseudo-register, 179
- threads, 10
 - active, listing, 46
 - blocked time, 458–473
 - blocking, 348
 - busy-spin loops, 347
 - call stacks, listing, 92
 - Cid, 522
 - CPU cores, bouncing between, 464
 - critical path of performance delays, 461
 - cyclic dependencies, 351
 - dangling, 337–338
 - freezing and unfreezing, 327
 - freezing and unfreezing in kernel mode, 262–265
 - freezing and unfreezing in user mode, 253–255
 - hash objects, assigning to, 328
 - impersonation, 353
 - information about, displaying, 521–522
 - kernel-mode and user-mode stacks, 383
 - leased, 102
 - listing, 149, 346
 - scheduling, 459–460
 - security contexts, 353
 - start/exit events, logging, 424
 - suspend count, 254
 - suspension, 15, 254–255, 262–265
 - suspension during garbage collection, 282
 - switching between, 511
 - user-mode and kernel-mode stacks, 297
 - wait states and single-stepping, 100–101
 - waits, viewing, 464
- !threads command, 149
- throw keyword, 89, 91
- thunking, 183
- timeout issues, 232–234
- time-travel debugging, 153
- tlbexp.exe, 28
- .tlist command, 59, 97
- tlist.exe, 131
- *.tmp extension, 243
- !token command, 59, 130–131, 152, 304, 355–357
 - n option, 355
- touch-capable applications, 20–21
- trace events. *See also* ETW events
 - logging, 391
- traces. *See also* ETW traces
 - collecting and analyzing, 153
 - merging event collections in, 418
 - save location, 398
- tracing. *See also* ETW (Event Tracing for Windows)
 - boot tracing, 449–454
 - circular ETW logging and, 421–422
 - CPU architecture and, 5
 - debugging with, 490–502
 - dynamically enabling, 312, 447
 - error code failures, 490–494
 - filtering, 443
 - handles, 302–303
 - in kernel-mode debugging, 76
 - lost events, 420
 - memory allocations, 311–312
 - with printf statements, 416
 - running processes, 399–400
 - starting with PerfView, 482
 - stopping with PerfView, 483
 - system call failures, 493–494
 - system internals, 494–502
 - time interval, 400
 - UAC elevation sequence, 495–502
- TracingFlags registry value, 477
- transport mediums for KD, 98
- trusted computing base (TCB), 10
- types
 - dumping, 48, 54, 80, 171, 238, 513–514
 - in public symbols files, 54
- type safety, 25

U

- UAC (User Account Control), 360–362, 495
 - consent.exe, 496–497
 - elevation requests, 461
 - elevation sequence, tracing, 495–502
 - improvements in Windows 7, 499
 - integrity levels and, 495

- ub command, 516
- !u command, 111
- ~u command, 254
- u command, 51, 97, 515
 - SOS version, 110–111
- uf command, 92, 212, 382, 515–516
- UI
 - message-loop handling, 367
 - unresponsive, 396–410
- UI events
 - asynchronous, 374–381
 - CancelKeyPress event, 374–375
 - console host process handling, 366–369, 373–374
 - main UI thread, 373
- UMDF (user-mode driver framework), 17–18
- UMDH tool
 - functionality, 311–312
 - memory leaks, investigating, 310–314
 - symbol resolution, 313–314
 - target process, running against instance of, 312
- unassembling, 51, 382
- unblocking, 467
- uncommitted page access violations, 271
- unfreezing threads, 253–255, 327
- unhandled exceptions, breaking on, 262
- unhandled verifier break assertions, 208
- Unicode strings, dumping values as, 52
- unloaded code
 - debugging, 341–342
 - !n command, 342
- unloaded DLL module access violations, 324, 340–343
- unloaded symbols, reloading, 341
- unmanaged-to-managed code transitions, 288–290
- unresolved breakpoints, 510
- unsafe C# code, 281–283
- unsigned int type, 89
- USB cables for kernel debugging, 68
- usbview.exe, 68
- USER32, 18
- user32.dll, 8, 18
- user access token, 353
- user authentication, 10
- user mode, 8–9
 - application processes, 10–13
 - crash-dump generation, 139–143
 - kernel-mode code, calling, 13–14, 15–16
 - kernel mode, invoking from, 15
 - kernel mode, transitioning to, 10
 - privileged tasks, performing, 10
 - system calls, 381–383
 - system processes, 9–10
- user-mode code
 - breakpoints, setting, 525
 - calling into kernel mode, 380–386
- user-mode crashes
 - blocking application from exiting, 125
 - debugging, 262
- user-mode debuggers
 - ALPC communication, 16
 - attaching to target, 220–222
 - debug port objects, attaching to, 237
 - !htrace command options, 302–303
 - security context, 220
 - security privileges, enabling, 493
 - simulating control by, 259–261
 - user context, 230
- user-mode debugging, 39–60
 - architecture of, 86–87
 - break-in sequence, 91–93
 - capabilities, 60
 - critical sections support, 347
 - data breakpoints, 164–165
 - debug events, 88–91
 - debugger functionality, 85–98
 - debugger prefix, 73
 - ending sessions, 518
 - exceptions handling, 88–91
 - quick start guide, 505–518
 - remote stubs, 119
 - starting sessions, 505
 - target process memory, inspecting, 86–87
 - Win32 debugging APIs, 87–88
 - with WinDbg, 39–47
- user-mode driver framework (UMDF), 17–18
- user-mode memory
 - code breakpoints, setting, 81–82
 - debugging leaks, 307–316
- user-mode processes
 - creation, debugging, 255–258
 - starting under debugger, 505
 - startup command line, 507
 - startup debugging, 255–258, 259–260
- user-mode stack, 297
- user-mode stack-trace database, 311
- snapshots of, 312–313
- user-mode system components, kernel debugging, 60

user provider events

- user provider events
 - logging during boot, 452–454
 - stack traces for, 434–436
 - viewing, 429–431
- user providers, 426–431
 - building and installing, 449
 - discovering, 427–429
 - enabling, 428, 443
 - friendly names and GUIDs, 427, 428
 - keyword bit masks, 443
 - manifest-based, 426–427
 - registering, 428
- user rights, 353
- user sessions, 9
 - user-mode processes, 10

V

- VBScript, 112
- vbscript.dll, 112
- verbose logging, 56
- verifier assertions, catching, 206
- verifier bits, OS-enabled, 216
- verifier break messages, 208
- verifier breaks
 - details, displaying, 276
 - information on, 215–216
 - resource-leak, 308–310
- verifier checks, 212
 - default set, 213
 - enabling, 207
- verifier.dll, 211–212
- verifier.exe, 213
- verifier flags, 214–217
- VerifierFlags value, 213
 - OS-enabled bits, 216
- !VerifyHeap command, 282
- vertarget command
 - CPU architecture information, 5
 - Windows version information, 6
- vfbasics.dll, 214
- View\Call Stack window, 514
- View\Locals menu action, 43, 513
- View\Memory window, 516
- View\Registers window, 43, 516
- view_trace.cmd script, 408, 438, 478
- View\Watch UI window, 513
- virtual address space
 - inspecting and modifying, 159
 - values, viewing, 268
- VirtualAlloc API, 268
- Virtual Clone Drive freeware, 35
- virtualization, 73–76
 - snapshots, 76
- virtual machines
 - kernel-mode debugging on, 73–76
 - restoring to snapshot, 76
 - testing, 75–76
- virtual memory
 - address space, 317
 - allocations, 268
 - leaks, 307
- Visual C++ (VC++)
 - catching bugs, 196–199
 - /GS compiler flag, 291–292
 - SAL annotations, 199–202
 - SEH exception support, 89
- Visual Studio, 29
 - Debug\Options And Settings dialog box, 132
 - downloading, 34
 - installing trial edition, 38
 - JIT debugging setup options, 132
 - out-of-process debugging support, 106
 - remote debugging in, 121–123
 - script debugging in, 114–115
 - script JIT debugging, 135–137
- Visual Studio 2010
 - automatic code analysis for, 196–197
 - code analysis warnings, 197–198
 - IntelliTrace, 153
 - Parallel Performance Analyzer, 458. *See also* PPA (Parallel Performance Analyzer)
 - profiling tools, 411
 - static code analysis, 195–206
 - Ultimate trial version, 38
 - wait analysis, 461–467
- Visual Studio debugger
 - analyzing crash dumps in, 150–151
 - attaching to target, 221
 - capabilities, 38–39
 - child debugging and, 238
 - DLL load events, breaking on, 225–227
 - entry point, 40
 - first-chance notifications, breaking on, 253
 - freezing and unfreezing threads, 254
 - installing, 37
 - JIT debugging, 132–137
 - managed-code debugging support, 104–105

- managed-code JIT debugging, 133–135
- native JIT debugging, 133
- .NET debugging, 103
- Process Debug Manager (PDM), 112
- remote debugging, 117
- scripts, attaching, 114–115
- as startup debugger, 230
- symbol decoration, 57–58
- vs. WinDbg, 38–39
- WinDbg, using with, 159
- Visual Studio Express edition, 38
- Visual Studio process instance (devenv.exe),
 - attaching WinDbg to, 235
- lvm command, 317
- VMMMap, memory-usage inspection, 474–475, 503
- VsJitDebugger.exe proxy process, 132

W

- wait analysis
 - call stacks, displaying, 465–466
 - in Visual Studio 2010, 461–467
 - in Xperf, 467–473
- WaitForMultipleObjects API, 333
- WaitForMultipleObjects(Ex) API, 170
- WaitForSingleObject API, 350
- WaitForSingleObject(Ex) API, 170
- waiting, 170–171
 - causality chains, 457
 - critical path of performance delays, 461
 - multiple conditions of, 460–461
 - unblocking, 467
- wait loops
 - attaching debugger during, 220–222
 - unblocking, 221
- wscpy function, 275, 294
- wscpy_s function, 275
- WDM (Windows Driver Model), 17
- WerFault.exe, 130
- WER (Windows Error Reporting)
 - automatic crash-dump generation, 139–143
 - JIT debugger invocation, 131
 - in Windows XP, 132
- wevtutil.exe tool, 427
- whoami.exe utility, 360
- wildcard character (*), 45
 - in bm command, 510
 - in k command, 512
 - in x command, 509
- Win32 API layer, 18–19
- Win32 APIs
 - calling conventions, 49–52
 - for debugging, 87–88
 - ETW logging, 445–449
 - failures, debugging, 245–247
 - SAL annotations, 202
 - system call transition, 380
- win32k.sys, 8, 80
- WinDbg (windbg.exe). *See also* Windows debuggers
 - Application Verifier and, 206
 - attaching to programs, 44
 - c command-line option, 177
 - child debugging, 234–245, 238
 - command prompt, 63
 - commands, capturing for scripts, 178
 - command window, 40
 - console applications, debugging, 370–373
 - crash dump analysis, 143–150
 - debugger command types, 59–60
 - debugging instance of, 93–97
 - DLL load events, breaking on, 223–225
 - elevated administrative command prompt, 63
 - ending debugging sessions, 43, 518, 526
 - entry point, advancing to, 40–41
 - function calls, listing, 47
 - as JIT debugger, 127–128
 - kernel-mode debugging quick start guide,
 - 519–526
 - live kernel debugging, 61–67
 - minidump generation, 152
 - multiple instances, 75
 - native flavor, 63
 - .NET applications, debugging, 107–112
 - noninvasive debugging, 159–161
 - pe command-line option, 161
 - pv command-line option, 160–161
 - remote debugging, 117–121
 - remote instance, 232
 - scripts, invoking, 177
 - SOS extension, 103, 106–112
 - source code window, 41
 - starting debugging sessions, 505, 519
 - symbols search path, setting, 45
 - user-mode debugging quick start guide,
 - 505–518
 - View\Call Stack window, 514
 - View\Memory window, 516
 - View\Registers window, 516
 - View\Watch UI window, 513

Windows 3.1

- Visual Studio debugger, using with, 159
- window arrangement, 41–43
- windows, docking, 41–42
- z command-line option, 142, 143
- Windows 3.1, 3
- Windows 7
 - UAC improvements, 499
 - Win32 DLL modules, 18
- Windows 7 Software Development Kit (SDK)
 - FxCop installation, 204
 - versions 7.0 and 7.1, 392
 - Windows debuggers, 34
 - Xperf, 391
- Windows 8 consumer preview, 395, 411
- Windows 32-bit on Windows 64-bit. *See* WOW64 debugging
- Windows 95/98/ME, 3
- Windows console subsystem, 365–380
 - Ctrl+C signal, handling, 374–380
 - printf function, 366–373
 - UI events, handling, 373–374
- Windows debuggers. *See also* WinDbg (windbg.exe)
 - capabilities, 38–39
 - commands, documentation, 58–60, 83
 - docking windows, 43
 - first-chance notifications, 91
 - Help files, 58–60, 83
 - installing, 34–37
 - int 3 instruction, inserting, 93
 - kernel-mode functionality, 98–102
 - local variables and parameters, listing values, 47–51
 - race conditions, simulating, 327–328
 - remote debugging, 117
 - scripting, 176–183
 - source-level debugging, lack of support, 52–53, 103, 106
 - symbol decoration, 57–58
 - symbols local cache, 55–56
 - system code, debugging, 44–47
 - use of, 33
 - user-mode debugging, 39–47
 - user-mode functionality, 85–98
 - vs. Visual Studio debugger, 38–39
 - x64 version, 36
 - x86 version, 36–37, 40
 - your own code, debugging, 39–43
- Windows Developer Preview Conference, 67
- windows, docking, 41–43
- Windows Driver Development Kit (DDK), 28–29
- Windows Driver Model (WDM), 17
- Windows Error Reporting (WER). *See* WER (Windows Error Reporting)
- Windows GUI process main UI thread, 373–374
- Windows Installer XML (WIX) declarative language, 240
- Windows logon process, 9–10
- Windows NT, 3
 - CPU architectures, 4
 - kernel, evolution of, 3–7
- Windows operating system
 - architecture, 7–16
 - boot sequence, tracing, 258
 - build flavors, 5–6
 - COM in, 20–21
 - communication mechanisms, 13–16
 - components, debugging, 33
 - CPU architectures, 4–5
 - developer interface, 16–28
 - developer tools, 28–29
 - ETW instrumentation, 422–431
 - evolution of, 3–7
 - exception dispatching sequence, 90
 - exception handling, 88–91
 - execution modes, 8–9
 - extending, 16–28
 - integrity levels, 359
 - managed-code debugging and, 105
 - message notifications, 86
 - prerelease milestones, 6
 - release history, 3–4
 - security model, 353–361
 - servicing updates, 7
 - verifier support, 209–214
- Windows PE (Portable Executable) format, 26
- Windows Performance Analyzer (WPA), 395, 411, 417
- Windows Performance Recorder (WPR), 395, 411, 417
- Windows Performance Toolkit (WPT). *See* WPT (Windows Performance Toolkit)
- Windows Phone SDK, 28
- Windows scripting hosts, 112
- Windows Server 2003, 3–4
- Windows Server 2008 R2 Hyper-V, 73
- Windows Server operating systems, 4
 - CPU architecture support, 5
- Windows service processes
 - killing, 492
 - listing, 181–183

- Windows services
 - handle leaks, 300
 - memory leaks, tracking, 310
 - starting, 10
 - startup debugging, 231–234
- Windows subsystem, 8
- Windows Vista, NT security model
 - improvements, 358–361
- Windows XP, 4
 - client/server code base, 3
 - CPU architectures, 5
- winlogon.exe, 9–10
- winnt.h Software Development Kit (SDK), 89
- winsrv!CreateCtrlThread function, 378
- __wmainCRTStartup function, 169, 336
- wmain function, 336
- workarounds, documentation of, 17
- worker processes
 - parent processes, 130
 - synchronizing with parent processes, 11
- worker threads
 - global variables in, 336–339
 - reference counting, 333–336
 - unloaded memory access violations, 340–343
 - waiting for, 333, 336
- wow64cpu.dll DLL, 183
- WOW64 debugging, 183–187
 - kernel-mode, 185–187
 - WOW64 environment, 183–184
- wow64.dll DLL, 183
- wow64exts.dll, 185
- wow64win.dll DLL, 183
- Wow6432Node registry key, 191
- WPA (Windows Performance Analyzer), 395, 411, 417
- wprintf function calls, SEH exceptions in, 198–199
- WPR (Windows Performance Recorder), 395, 411, 417
- WPT (Windows Performance Toolkit), 172
 - Help file, 395
 - installing, 392–395
 - native version, 393
 - Windows 8 version, 395
- write-behind optimization, 403
- .write_cmd_hist command, 178
- write-data breakpoints, 294
- WriteFile API, SAL annotations, 202
- WriteProcessMemory API
 - arguments, checking, 96
 - parameters, 96–97

- wscript.exe, 112
 - //X option, 114

X

- x64 debuggers, 184
- x64 processes
 - debugging, 78
 - virtual address space, 317
- x64 processors, 4–5, 183
 - Windows debugger version, 36
- x64 Windows
 - syscall and sysret instructions, 380
 - x64 WPT, 393
- x86 applications. *See also* WOW64 debugging
 - x64 support for, 183
- x86 debuggers, 184
- x86 processes
 - debugging, 78
 - virtual memory address space, 317
- x86 processors, 4–5
 - Windows debugger version, 36–37
- x86 registers, 183
- x86 Windows
 - frame pointer omission optimization, 440
 - script for, 257
 - sysenter and sysexit instructions, 380
 - x86 WPT, 393
- Xbox development kit (XDK), 28
- x command, 45, 57, 181, 222, 260, 261, 509
 - //X command-line option, 135
- XmlTextReader class, 487
- Xperf (xperf.exe), 391–413, 499
 - a command-line option, 444
 - acquiring, 391–395
 - allocation types, 480
 - BootTrace option, 450
 - BufferSize, –MinBuffers, and –MaxBuffers
 - command-line options, 420
 - Column Chooser flyout, 406–407
 - command-line options and features, 395
 - CPU sample profiling, 502
 - CPU Sampling By Process graph, 405
 - CPU sampling profile graph, 404–405
 - CPU Sampling Summary Table window, 406–407
 - CPU Scheduling graph, 468–469
 - d command-line option, 399, 411, 418, 438
 - Disk Utilization graph, 402–404
 - dumper action, 444

Xperf (xperf.exe)

- Xperf (xperf.exe), *continued*
 - ETW heap traces, capturing, 476–480
 - ETW logging, starting, 418
 - ETW sessions, configuring, 419–422
 - ETW traces, analyzing, 399–410, 444–445
 - ETW traces, collecting, 397–398
 - FileMode option, 421
 - frame list flyout, 402
 - Generic Events graph, 426, 429–430, 436, 500
 - gold bar, 407, 409
 - graphs and data views, selecting, 402–405
 - graphs, overlaying, 431
 - heap memory leak investigations, 503
 - heap option, 476
 - Heap Outstanding Allocation Size graph, 479
 - Heap Total Allocation Size graph, 478
 - incomplete stack-trace events, 439–441
 - investigating with, 396–410
 - investigation strategy, 397
 - kernel flag combinations, 422
 - kernel provider events, viewing, 425–426
 - in live production environments, 411
 - location, 394–395
 - loggers, 417
 - loggers command-line option, 421
 - lost events warning, 420
 - main UI window, 399
 - mark events, 430–431
 - missing stack-trace events, 436–437
 - m option, 430
 - _NT_SYMCACHE_PATH environment variable, 408
 - on command-line option, 418
 - performance delays, investigating, 502
 - Process Lifetimes graph, 399–400, 423, 451
 - Process Summary Table, 400–401
 - providers option, 427
 - scope of analysis, 397
 - search path, 395
 - session configuration options, 419, 422
 - Stack Counts By Type graph, 433, 494
 - Stack Counts By Type summary table, 498
 - Stack Counts Summary Table, 454
 - :::‘stack’ string, 434, 443
 - stackwalk option, 432, 434
 - start command-line option, 418
 - Start/End event pair, viewing, 448
 - stop command-line option, 418
 - strengths and limitations, 411
 - Summary Table view, 435
 - symbols, loading, 405–406, 433–435, 438
 - symcache action, 444
 - system configuration information, 401
 - system-level tracing, 503
 - time interval of trace, 400
 - trace files, viewing, 399
 - transition to WPR/WPA, 395
 - unresolved stack-trace events, 437–439
 - user provider events, viewing, 429–431
 - user providers, 426–431
 - wait analysis, 467–473

Y

- your own code
 - debugging, 39–43
 - parameters and locals, listing, 48–49

Z

- z command-line option, 57, 142, 143

About the Author



TARIK SOULAMI, a principal development lead on the Windows Fundamentals Team, has more than 10 years of experience designing and developing system-level software at Microsoft. Before joining the Windows team, Tarik spent several years on the Common Language Runtime (CLR) Team, where he helped shape early releases of the Microsoft .NET framework.