


Windows via C/C++ Fifth Edition

Windows核心编程

(第5版·英文版)

[美] Jeffrey Richter 著
[法] Christophe Nasarre

- Windows程序设计巅峰之作
- 新版涵盖Windows Vista和Windows Server 2008最新核心特性
- Windows程序员迈向卓越的必由之路



TURING

图灵程序设计丛书

微软技术系列

Windows via C/C++

Fifth Edition

Windows核心编程

(第5版·英文版)

[美] Jeffrey Richter

[法] Christophe Nasarre 著

Jeffrey Richter

人民邮电出版社
北京

图书在版编目 (CIP) 数据

Windows 核心编程: 第 5 版: 英文 / (美) 里克特 (Richter, J.), (法) 纳萨尔 (Nasarre, C.) 著; 一北京: 人民邮电出版社, 2008.8
(图灵程序设计丛书)
ISBN 978-7-115-18212-8

I. W… II. ①里…②纳… III. 窗口软件, Windows—程序设计—英文 IV. TP316.7

中国版本图书馆 CIP 数据核字 (2008) 第 077689 号

内 容 提 要

本书是 Windows 程序设计领域的名著, 涵盖了 Windows 的最新版本 Vista 以及 Windows XP 的最新内容。书中全面深入地介绍了 Windows 的各种基本要素, 如进程、线程池、虚拟内存、DLL、设备 I/O 和 SEH 等, 并列出了大量应用程序, 精辟地分析了要素的使用方法。

本书适于各层次 Windows 编程人员阅读。

图灵程序设计丛书

Windows 核心编程 (第 5 版·英文版)

-
- ◆ 著 [美] Jeffrey Richter [法] Christophe Nasarre
责任编辑 傅志红
 - ◆ 人民邮电出版社出版发行 北京市崇文区夕照寺街 14 号
邮编 100061 电子函件 315@ptpress.com.cn
网址 <http://www.ptpress.com.cn>
三河市潮河印业有限公司印刷
 - ◆ 开本: 800×1000 1/16
印张: 52.75
字数: 1013 千字 2008 年 8 月第 1 版
印数: 1-3 000 册 2008 年 8 月河北第 1 次印刷

著作权合同登记号 图字: 01-2008-1720 号

ISBN 978-7-115-18212-8/TP

定价: 128.00 元

读者服务热线: (010)88593802 印装质量热线: (010)67129223
反盗版热线: (010)67171154

版 权 声 明

Original edition, entitled *Windows via C/C++ fifth, Edition* by Jeffrey Richter and Christophe Nasarre, ISBN 9780735624245 published by Microsoft Press in 2008.

This reprint edition is published with the permission of the Syndicate of the Microsoft Press.

Copyright © 2008 by Jeffrey Richter and Christophe Nasarre.

THIS EDITION IS LICENSED FOR DISTRIBUTION AND SALE IN THE PEOPLE'S REPUBLIC OF CHINA ONLY, EXCLUDING HONG KONG, MACAO AND TAIWAN, AND MAY NOT BE DISTRIBUTED AND SOLD ELSEWHERE.

本书原版由微软出版社出版。

本书英文影印版由微软出版社授权人民邮电出版社独家出版。未经出版者书面许可，不得以任何方式复制或抄袭本书内容。

此版本仅限在中华人民共和国境内（不包括中国香港、澳门特别行政区及中国台湾地区）销售发行。

版权所有，侵权必究。

Dedication

*To Kristin, words cannot express how I feel about our life together.
I cherish our family and all our adventures.
I'm filled each day with love for you.*

*To Aidan, you have been an inspiration to me and have taught me
to play and have fun. Watching you grow up has been so rewarding
and enjoyable for me. I feel lucky to be able to
partake in your life; it has made me a better person.*

*To My New Baby Boy (shipping Q1 2008),
you have been wanted for so long it's hard to believe
that you're almost here. You bring completeness and balance
to our family. I look forward to playing with you,
learning who you are, and enjoying our time together.*

– Jeffrey Richter

To my wife Florence, au moins cette fois c'est écrit: je t'aime Flo.

*To my parents who cannot believe that learning English
with Dungeons & Dragons rules could have been so efficient.*

– Christophe Nasarre

Introduction

Microsoft Windows is a complex operating system. It offers so many features and does so much that it's impossible for any one person to fully understand the entire system. This complexity also makes it difficult for someone to decide where to start concentrating the learning effort. Well, I always like to start at the lowest level by gaining a solid understanding of the system's basic building blocks. Once you understand the basics, it's easy to incrementally add any higher-level aspects of the system to your knowledge. So this book focuses on Windows' basic building blocks and the fundamental concepts that you must know when architecting and implementing software targeting the Windows operating system. In short, this book teaches the reader about various Windows features and how to access them via the C and C++ programming languages.

Although this book does not cover some Windows concepts—such as the Component Object Model (COM)—COM is built on top of basic building blocks such as processes, threads, memory management, DLLs, thread local storage, Unicode, and so on. If you know these basic building blocks, understanding COM is just a matter of understanding how the building blocks are used. I have great sympathy for people who attempt to jump ahead in learning COM's architecture. They have a long road ahead and are bound to have gaping holes in their knowledge, which is bound to negatively affect their code and their software development schedules.

The Microsoft .NET Framework's common language runtime (CLR) is another technology not specifically addressed in this book. (However, it is addressed in my other book: *CLR via C#*, Jeffrey Richter, Microsoft Press, 2006). However, the CLR is implemented as a COM object in a dynamic-link library (DLL) that loads in a process and uses threads to execute code that manipulates Unicode strings that are managed in memory. So again, the basic building blocks presented in this book will help developers writing managed code. In addition, by way of the CLR's Platform Invocation (P/Invoke) technology, you can call into the various Windows' APIs presented throughout this book.

So that's what this book is all about: the basic Windows building blocks that every Windows developer (at least in my opinion) should be intimately aware of. As each block is discussed, I also describe how the system uses these blocks and how your own applications can best take advantage of these blocks. In many chapters, I show you how to create building blocks of your own. These building blocks, typically implemented as generic functions or C++ classes, group a set of Windows building blocks together to create a whole that is much greater than the sum of its parts.

64-Bit Windows

Microsoft has been shipping 32-bit versions of Windows that support the x86 CPU architecture for many years. Today, Microsoft also offers 64-bit versions of Windows that support the x64 and IA-64 CPU architectures. Machines based on these 64-bit CPU architectures are fast gaining acceptance. In fact, in the very near future, it is expected that all desktop and server machines will contain 64-bit CPUs. Because of this, Microsoft has stated that Windows Server 2008 will be the last 32-bit version of Windows ever! For developers, now is the time to focus on making sure your applications run correctly on 64-bit Windows. To this end, this book includes solid coverage of what you need to know to have your applications run on 64-bit Windows (as well as 32-bit Windows).

The biggest advantage your application gets from a 64-bit address space is the ability to easily manipulate large amounts of data, because your process is no longer constrained to a 2-GB usable address space. Even if your application doesn't need all this address space, Windows itself takes advantage of the significantly larger address space (about 8 terabytes), allowing it to run faster.

Here is a quick look at what you need to know about 64-bit Windows:

- The 64-bit Windows kernel is a port of the 32-bit Windows kernel. This means that all the details and intricacies that you've learned about 32-bit Windows still apply in the 64-bit world. In fact, Microsoft has modified the 32-bit Windows source code so that it can be compiled to produce a 32-bit or a 64-bit system. They have just one source-code base, so new features and bug fixes are simultaneously applied to both systems.
- Because the kernels use the same code and underlying concepts, the Windows API is identical on both platforms. This means that you do not have to redesign or reimplement your application to work on 64-bit Windows. You can simply make slight modifications to your source code and then rebuild.
- For backward compatibility, 64-bit Windows can execute 32-bit applications. However, your application's performance will improve if the application is built as a true 64-bit application.
- Because it is so easy to port 32-bit code, there are already device drivers, tools, and applications available for 64-bit Windows. Unfortunately, Visual Studio is a native 32-bit application and Microsoft seems to be in no hurry to port it to be a native 64-bit application. However, the good news is that 32-bit Visual Studio does run quite well on 64-bit Windows; it just has a limited address space for its own data structures. And Visual Studio does allow you to debug a 64-bit application.
- There is little new for you to learn. You'll be happy to know that most data types remain 32 bits wide. These include **ints**, **DWORDs**, **LONGs**, **BOOLs**, and so on. In fact, you mostly just need to worry about pointers and handles, since they are now 64-bit values.

Because Microsoft offers so much information on how to modify your existing source code to be 64-bit ready, I will not go into those details in this book. However, I thought about 64-bit Windows as I wrote each chapter. Where appropriate, I have included information specific to 64-bit Windows. In addition, I have compiled and tested all the sample applications in this book for 64-bit Windows. So, if you follow the sample applications in this book and do as I've done, you should have no trouble creating a single source-code base that you can easily compile for 32-bit or 64-bit Windows.

What's New in the Fifth Edition

In the past, this book has been titled *Advanced Windows NT*, *Advanced Windows*, and *Programming Applications for Microsoft Windows*. In keeping with tradition, this edition of the book has gotten a new title: *Windows via C/C++*. This new title indicates that the book is for C and C++ programmers wanting to understand Windows. This new edition covers more than 170 new functions and Windows features that have been introduced in Windows XP, Windows Vista, and Windows Server 2008.

Some chapters have been completely rewritten—such as Chapter 11, which explains how the new thread pool API should be used. Existing chapters have been greatly enhanced to present new features. For example, Chapter 4 now includes coverage of User Account Control and Chapter 8 now covers new synchronization mechanisms (Interlocked Singly-Linked List, Slim Reader-Writer Locks, and condition variables).

I also give much more coverage of how the C/C++ run-time library interacts with the operating system—particularly on enhancing security as well as exception handling. Last but not least, two new chapters have been added to explain how I/O operations work and to dig into the new Windows Error Reporting system that changes the way you must think about application error reporting and application recovery.

In addition to the new organization and greater depth, I added a ton of new content. Here is a partial list of enhancements made for this edition:

New Windows Vista and Windows Server 2008 features Of course, the book would not be a true revision unless it covered new features offered in Windows XP, Windows Vista, Windows Server 2008, and the C/C++ run-time library. This edition has new information on the secure string functions, the kernel object changes (such as namespaces and boundary descriptors), thread and process attribute lists, thread and I/O priority scheduling, synchronous I/O cancellation, vectored exception handling, and more.

64-bit Windows support The text addresses 64-bit Windows-specific issues; all sample applications have been built and tested on 64-bit Windows.

Use of C++ The sample applications use C++ and require fewer lines of code, and their logic is easier to follow and understand.

Reusable code Whenever possible, I created the source code to be generic and reusable. This should allow you to take individual functions or entire C++ classes and drop them into your own applications with little or no modification. The use of C++ made reusability much easier.

The ProcessInfo utility This particular sample application from the earlier editions has been enhanced to show the process owner, command line, and UAC-related details.

The LockCop utility This sample application is new. It shows which processes are running on the system. Once you select a process, this utility lists the threads of the process and, for each, on which kind of synchronization mechanism it is blocked—with deadlocks explicitly pointed out.

API hooking I present updated C++ classes that make it trivial to hook APIs in one or all modules of a process. My code even traps run-time calls to **LoadLibrary** and **GetProcAddress** so that your API hooks are enforced.

Structured exception handling improvements I have rewritten and reorganized much of the structured exception handling material. I have more information on unhandled exceptions, and I've added coverage on customizing Windows Error Reporting to fulfill your needs.

Code Samples and System Requirements

The sample applications presented throughout this book can be downloaded from the book's companion content Web page at

<http://www.Wintellect.com/Books.aspx>

To build the applications, you'll need Visual Studio 2005 (or later), the Microsoft Platform SDK for Windows Vista and Windows Server 2008 (which comes with some versions of Visual Studio). In addition, to run the applications, you'll need a computer (or virtual machine) with Windows Vista (or later) installed.

Support for This Book

Every effort has been made to ensure the accuracy of this book and the companion content. As corrections or changes are collected, they will be added to an Errata document downloadable at the following Web site:

<http://www.Wintellect.com/Books.aspx>

Questions and Comments

If you have comments, questions, or ideas regarding the book or the companion content, or questions that are not answered by visiting the site just mentioned, please send them to Microsoft Press via e-mail to

mspinput@microsoft.com

Or via postal mail to

Microsoft Press
Attn: *Windows via C/C++* Editor
One Microsoft Way
Redmond, WA 98052-6399

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

Acknowledgments

We could not have written this book without the help and technical assistance of several people. In particular, we'd like to thank the following people:

Jeffrey's Family

Jeffrey would like to thank Kristin (his wife) and Aidan (his son) for their never ending love and support.

Christophe's Family

Christophe would not have been able to write the fifth edition of this book without the love and support of Florence (his wife), the never ending curiosity of Celia (his daughter), and the purring sneak attacks of Canelle and Nougat (his cats). Now, I don't have any good excuse to not take care of you!

Technical Assistance

For writing a book like this one, personal research is not enough. We owe great thanks to various Microsoft employees who helped us. Specifically, we'd like to thank Arun Kishan, who was able to either instantly answer weird and complicated questions or find the right person on the Windows team to provide more detailed explanations. We would also like to thank Kinshuman Kinshumann, Stephan Doll, Wedson Almeida Filho, Eric Li, Jean-Yves Pouban, Sandeep Ranade, Alan Chan, Ale Contenti, Kang Su Gatlin, Kai Hsu, Mehmet Iyigun, Ken Jung, Pavel Lebedynskiy, Paul Sliwicz, and Landy Wang. In addition, there are those who listened to questions posted on Microsoft internal forums and shared their extensive knowledge, such as Raymond Chen, Sungook Chue, Chris Corio, Larry Osterman, Richard Russell, Mark Russinovich, Mike Sheldon, Damien Watkins, and Junfeng Zhang. Last but not least, we would like to warmly thank John "Bugslayer" Robbins and Kenny Kerr who were kind enough to provide great feedback on chapters of this book.

Microsoft Press Editorial Team

We would like to thank Ben Ryan (acquisitions editor) for trusting a crazy French guy like Christophe, managers Lynn Finnel and Curtis Philips for their patience, Scott Seely for his constant search for technical accuracy, Roger LeBlanc for his talent in transforming Christophe's French-like English into something understandable, and Andrea Fox for her meticulous proofreading. In addition to the Redmond team, Joyanta Sen spent a lot of his personal time supporting us.

Mutual Admiration

Christophe sincerely thanks Jeffrey Richter for trusting him not to spoil the fifth edition of Jeff's book.

Jeffrey also thanks Christophe for his tireless efforts in researching, reorganizing, rewriting, and reworking the content in an attempt to reach Jeff's idea of perfection.

Table of Contents

Part I **Required Reading**

1	Error Handling	3
	Defining Your Own Error Codes	7
	The ErrorShow Sample Application	7
2	Working with Characters and Strings	11
	Character Encodings	12
	ANSI and Unicode Character and String Data Types	13
	Unicode and ANSI Functions in Windows	15
	Unicode and ANSI Functions in the C Run-Time Library	17
	Secure String Functions in the C Run-Time Library	18
	Introducing the New Secure String Functions	19
	How to Get More Control When Performing String Operations	22
	Windows String Functions	24
	Why You Should Use Unicode	26
	How We Recommend Working with Characters and Strings	26
	Translating Strings Between Unicode and ANSI	27
	Exporting ANSI and Unicode DLL Functions	29

2 Table of Contents

Determining If Text Is ANSI or Unicode 31

3 Kernel Objects 33

What Is a Kernel Object? 33

Usage Counting 35

Security 35

A Process' Kernel Object Handle Table 37

Creating a Kernel Object 38

Closing a Kernel Object 39

Sharing Kernel Objects Across Process Boundaries 43

Using Object Handle Inheritance 43

Naming Objects 48

Duplicating Object Handles 60

Part II Getting Work Done

4 Processes 67

Writing Your First Windows Application 68

A Process Instance Handle 73

The *CreateProcess* Function 89

pszApplicationName and *pszCommandLine* 89

Terminating a Process 104

The Primary Thread's Entry-Point Function Returns 104

The *ExitProcess* Function 105

The *TerminateProcess* Function 106

When All the Threads in the Process Die 107

When a Process Terminates 107

Child Processes 108

Running Detached Child Processes 110

When Administrator Runs as a Standard User 110

Elevating a Process Automatically 113

Elevating a Process by Hand 115

What Is the Current Privileges Context? 117

Enumerating the Processes Running in the System 118

5	Jobs	125
	Placing Restrictions on a Job's Processes	129
	Placing a Process in a Job	136
	Terminating All Processes in a Job	136
	Querying Job Statistics	137
	Job Notifications	140
	The Job Lab Sample Application	143
6	Thread Basics	145
	When to Create a Thread	146
	When Not to Create a Thread	148
	Writing Your First Thread Function	149
	The <i>CreateThread</i> Function	150
	<i>psa</i>	151
	<i>cbStackSize</i>	151
	<i>pfnStartAddr</i> and <i>pvParam</i>	152
	<i>dwCreateFlags</i>	153
	<i>pdwThreadId</i>	153
	Terminating a Thread	153
	The Thread Function Returns	154
	The <i>ExitThread</i> Function	154
	The <i>TerminateThread</i> Function	154
	When a Process Terminates	155
	When a Thread Terminates	155
	Some Thread Internals	156
	C/C++ Run-Time Library Considerations	159
	Oops—I Called <i>CreateThread</i> Instead of <i>_beginthreadex</i> by Mistake	168
	C/C++ Run-Time Library Functions That You Should Never Call	168
	Gaining a Sense of One's Own Identity	169
	Converting a Pseudohandle to a Real Handle	170

7	Thread Scheduling, Priorities, and Affinities	173
	Suspending and Resuming a Thread	175
	Suspending and Resuming a Process	176
	Sleeping	177
	Switching to Another Thread	178
	Switching to Another Thread on a Hyper-Threaded CPU	178
	A Thread's Execution Times	179
	Putting the <i>CONTEXT</i> in Context	183
	Thread Priorities	187
	An Abstract View of Priorities	188
	Programming Priorities	191
	Dynamically Boosting Thread Priority Levels	194
	Tweaking the Scheduler for the Foreground Process	195
	Scheduling I/O Request Priorities	196
	The Scheduling Lab Sample Application	197
	Affinities	203
8	Thread Synchronization in User Mode	207
	Atomic Access: The Interlocked Family of Functions	208
	Cache Lines	214
	Advanced Thread Synchronization	215
	A Technique to Avoid	216
	Critical Sections	217
	Critical Sections: The Fine Print	219
	Critical Sections and Spinlocks	222
	Critical Sections and Error Handling	223
	Slim Reader-Writer Locks	224
	Condition Variables	227
	The Queue Sample Application	228
	Useful Tips and Techniques	238

9	Thread Synchronization with Kernel Objects	241
	Wait Functions	243
	Successful Wait Side Effects	246
	Event Kernel Objects	247
	The Handshake Sample Application	252
	Waitable Timer Kernel Objects	256
	Having Waitable Timers Queue APC Entries	260
	Timer Loose Ends	261
	Semaphore Kernel Objects	262
	Mutex Kernel Objects	265
	Abandonment Issues	267
	Mutexes vs. Critical Sections	267
	The Queue Sample Application	268
	A Handy Thread Synchronization Object Chart	276
	Other Thread Synchronization Functions	277
	Asynchronous Device I/O	277
	<i>WaitForInputIdle</i>	278
	<i>MsgWaitForMultipleObjects(Ex)</i>	278
	<i>WaitForDebugEvent</i>	279
	<i>SignalObjectAndWait</i>	279
	Detecting Deadlocks with the Wait Chain Traversal API	281
10	Synchronous and Asynchronous Device I/O	289
	Opening and Closing Devices	290
	A Detailed Look at <i>CreateFile</i>	292
	Working with File Devices	299
	Getting a File's Size	299
	Positioning a File Pointer	300
	Setting the End of a File	302
	Performing Synchronous Device I/O	302
	Flushing Data to the Device	303
	Synchronous I/O Cancellation	303

6 Table of Contents

- Basics of Asynchronous Device I/O 305
 - The *OVERLAPPED* Structure 306
 - Asynchronous Device I/O Caveats 307
 - Canceling Queued Device I/O Requests 309
- Receiving Completed I/O Request Notifications 310
 - Signaling a Device Kernel Object 311
 - Signaling an Event Kernel Object 312
 - Alertable I/O 315
 - I/O Completion Ports 320

11 The Windows Thread Pool. 339

- Scenario 1: Call a Function Asynchronously 340
 - Explicitly Controlling a Work Item 340
 - The Batch Sample Application 342
- Scenario 2: Call a Function at a Timed Interval 346
 - The Timed Message Box Sample Application 348
- Scenario 3: Call a Function When a Single Kernel Object Becomes Signaled. 351
- Scenario 4: Call a Function When Asynchronous I/O Requests Complete 353
- Callback Termination Actions 355
 - Customized Thread Pools 356
 - Gracefully Destroying a Thread Pool: Cleanup Groups. 358

12 Fibers. 361

- Working with Fibers 361
 - The Counter Sample Application 365

Part III Memory Management

13 Windows Memory Architecture 371

- A Process' Virtual Address Space 371
- How a Virtual Address Space Is Partitioned 372
 - Null-Pointer Assignment Partition 372
 - User-Mode Partition 373
 - Kernel-Mode Partition 375

Regions in an Address Space	375
Committing Physical Storage Within a Region	376
Physical Storage and the Paging File	377
Physical Storage Not Maintained in the Paging File	379
Protection Attributes	381
Copy-on-Write Access	382
Special Access Protection Attribute Flags	382
Bringing It All Home	383
Inside the Regions	388
The Importance of Data Alignment	391
14 Exploring Virtual Memory	395
System Information	395
The System Information Sample Application	398
Virtual Memory Status	404
Memory Management on NUMA Machines	405
The Virtual Memory Status Sample Application	406
Determining the State of an Address Space	408
The <i>VMQuery</i> Function	410
The Virtual Memory Map Sample Application	415
15 Using Virtual Memory in Your Own Applications	419
Reserving a Region in an Address Space	419
Committing Storage in a Reserved Region	421
Reserving a Region and Committing Storage Simultaneously	422
When to Commit Physical Storage	424
Decommitting Physical Storage and Releasing a Region	426
When to Decommit Physical Storage	426
The Virtual Memory Allocation Sample Application	427
Changing Protection Attributes	434
Resetting the Contents of Physical Storage	435
The MemReset Sample Application	437
Address Windowing Extensions	439
The AWE Sample Application	442