TURING 图灵程序设计丛书

Microsoft

Windows Internals

Part 1 Sixth Edition

深入解析Windows操作系统

卷1（英文版·第6版）

[美] Mark Russinovich
David Solomon
[加] Alex Ionescu
著

- 微软官方权威著作最新版
- 深入剖析Windows技术内幕
- 大幅更新，涵盖Windows内核最新特性

人民邮电出版社
POSTS & TELECOM PRESS

TURING 图灵程序设计丛

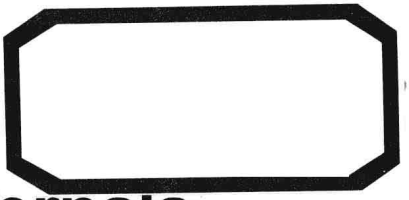

Windows Internals

Part 1 Sixth Edition

深入解析Windows操作系统

卷1（英文版·第6版）

[美] Mark Russinovich
David Solomon
[加] Alex Ionescu
著

人民邮电出版社
北京

图书在版编目（CIP）数据

深入解析Windows操作系统 : 第6版. 卷1 = Windows Internals:Sixth Edition.Part 1 : 英文 / (美) 拉希诺维奇 (Russinovich,M.) , (美) 所罗门 (Solomon,D.), (加) 艾欧内斯库 (Ionescu,A.) 著. -- 北京 : 人民邮电出版社, 2012.9
（图灵程序设计丛书）
ISBN 978-7-115-29090-8

Ⅰ. ①深… Ⅱ. ①拉… ②所… ③艾… Ⅲ. ①Windows操作系统—英文 Ⅳ. ①TP316.7

中国版本图书馆CIP数据核字(2012)第188455号

内 容 提 要

本书是操作系统内核专家 Russinovich 等人的 Windows 操作系统原理的最新版著作，针对 Windows 7 和 Windows Server 2008 R2 进行了全面的更新，主要讲述 Windows 的底层关键机制、Windows 的核心组件（包括进程 / 线程 / 作业、安全性、I/O 系统、存储管理、内存管理、缓存管理、文件系统和网络），并分析了启动进程、关机进程以及缓存转储。书中提供了许多实例，读者可以借此更好地理解 Windows 的内部行为。

本书内容丰富，信息全面，适合众多 Windows 平台开发人员、系统管理员阅读。

图灵程序设计丛书

深入解析Windows操作系统，卷1（英文版·第6版）

◆ 著 [美] Mark Russinovich David Solomon
[加] Alex Ionescu

责任编辑 朱 巍

◆ 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 电子邮件 315@ptpress.com.cn
网址 http://www.ptpress.com.cn
北京艺辉印刷有限公司印刷

◆ 开本：800×1000 1/16
印张：46.5
字数：744千字 2012年9月第1版
印数：1-3 000册 2012年9月北京第1次印刷

著作权合同登记号 图字：01-2012-4474 号

ISBN 978-7-115-29090-8

定价：99.00元

读者服务热线：(010)51095186转604 印装质量热线：(010)67129223

反盗版热线：(010)67171154

版权声明

Introduction

Windows Internals, Sixth Edition is intended for advanced computer professionals (both developers and system administrators) who want to understand how the core components of the Microsoft Windows 7 and Windows Server 2008 R2 operating systems work internally. With this knowledge, developers can better comprehend the rationale behind design choices when building applications specific to the Windows platform. Such knowledge can also help developers debug complex problems. System administrators can benefit from this information as well, because understanding how the operating system works "under the covers" facilitates understanding the performance behavior of the system and makes troubleshooting system problems much easier when things go wrong. After reading this book, you should have a better understanding of how Windows works and why it behaves as it does.

Structure of the Book

For the first time, *Windows Internals* has been divided into two parts. Updating the book for each release of Windows takes considerable time so producing it in two parts allows us to publish the first part earlier.

This book, Part 1, begins with two chapters that define key concepts, introduce the tools used in the book, and describe the overall system architecture and components. The next two chapters present key underlying system and management mechanisms. Part 1 wraps up by covering three core components of the operating system: processes, threads, and jobs; security; and networking.

Part 2, which will be available separately in fall 2012, covers the remaining core subsystems: I/O, storage, memory management, the cache manager, and file systems. Part 2 concludes with a description of the startup and shutdown processes and a description of crash-dump analysis.

History of the Book

This is the sixth edition of a book that was originally called *Inside Windows NT* (Microsoft Press, 1992), written by Helen Custer (prior to the initial release of Microsoft Windows NT 3.1). *Inside Windows NT* was the first book ever published about Windows NT and provided key insights into the architecture and design of the system. *Inside Windows NT, Second Edition* (Microsoft Press, 1998) was written by David Solomon. It updated the original book to cover Windows NT 4.0 and had a greatly increased level of technical depth.

Inside Windows 2000, Third Edition (Microsoft Press, 2000) was authored by David Solomon and Mark Russinovich. It added many new topics, such as startup and shutdown, service internals, registry internals, file-system drivers, and networking. It also covered kernel changes in Windows 2000, such as the Windows Driver Model (WDM), Plug and Play, power management, Windows Management Instrumentation (WMI), encryption, the job object, and Terminal Services. *Windows Internals, Fourth Edition* was the Windows XP and Windows Server 2003 update and added more content focused on helping IT professionals make use of their knowledge of Windows internals, such as using key tools from Windows Sysinternals (*www.microsoft.com/technet/sysinternals*) and analyzing crash dumps. *Windows Internals, Fifth Edition* was the update for Windows Vista and Windows Server 2008. New content included the image loader, user-mode debugging facility, and Hyper-V.

Sixth Edition Changes

This latest edition has been updated to cover the kernel changes made in Windows 7 and Windows Server 2008 R2. Hands-on experiments have been updated to reflect changes in tools.

Hands-on Experiments

Even without access to the Windows source code, you can glean much about Windows internals from tools such as the kernel debugger and tools from Sysinternals and Winsider Seminars & Solutions. When a tool can be used to expose or demonstrate some aspect of the internal behavior of Windows, the steps for trying the tool yourself are listed in "EXPERIMENT" boxes. These appear throughout the book, and we encourage you to try these as you're reading—seeing visible proof of how Windows works internally will make much more of an impression on you than just reading about it will.

Topics Not Covered

Windows is a large and complex operating system. This book doesn't cover everything relevant to Windows internals but instead focuses on the base system components. For example, this book doesn't describe COM+, the Windows distributed object-oriented programming infrastructure, or the Microsoft .NET Framework, the foundation of managed code applications.

Because this is an internals book and not a user, programming, or system administration book, it doesn't describe how to use, program, or configure Windows.

A Warning and a Caveat

Because this book describes undocumented behavior of the internal architecture and the operation of the Windows operating system (such as internal kernel structures and functions), this content is subject to change between releases. (External interfaces, such as the Windows API, are not subject to incompatible changes.)

By "subject to change," we don't necessarily mean that details described in this book will change between releases, but you can't count on them not changing. Any software that uses these undocumented interfaces might not work on future releases of Windows. Even worse, software that runs in kernel mode (such as device drivers) and uses these undocumented interfaces might experience a system crash when running on a newer release of Windows.

Acknowledgments

First, thanks to Jamie Hanrahan and Brian Catlin of Azius, LLC for joining us on this project—the book would not have been finished without their help. They did the bulk of the updates on the "Security" and "Networking" chapters and contributed to the update of the "Management Mechanisms" and "Processes and Threads" chapters. Azius provides Windows-internals and device-driver training. See *www.azius.com* for more information.

We want to recognize Alex Ionescu, who for this edition is a full coauthor. This is a reflection of Alex's extensive work on the fifth edition, as well as his continuing work on this edition.

Thanks to Eric Traut and Jon DeVaan for continuing to allow David Solomon access to the Windows source code for his work on this book as well as continued development of his Windows Internals courses.

Three key reviewers were not acknowledged for their review and contributions to the fifth edition: Arun Kishan, Landy Wang, and Aaron Margosis—thanks again to them! And thanks again to Arun and Landy for their detailed review and helpful input for this edition.

This book wouldn't contain the depth of technical detail or the level of accuracy it has without the review, input, and support of key members of the Microsoft Windows development team. Therefore, we want to thank the following people, who provided technical review and input to the book:

- Greg Cottingham
- Joe Hamburg
- Jeff Lambert
- Pavel Lebedynskiy
- Joseph East
- Adi Oltean
- Alexey Pakhunov
- Valerie See

For the "Networking" chapter, a special thanks to Gianluigi Nusca and Tom Jolly, who really went beyond the call of duty: Gianluigi for his extraordinary help with the BranchCache material and the amount of suggestions (and many paragraphs of material he wrote), and Tom Jolly not only for his own review and suggestions (which were excellent), but for getting many other developers to assist with the review. Here are all those who reviewed and contributed to the "Networking" chapter:

- Roopesh Battepati
- Molly Brown
- Greg Cottingham
- Dotan Elharrar
- Eric Hanson
- Tom Jolly

- Manoj Kadam
- Greg Kramer
- David Kruse
- Jeff Lambert
- Darene Lewis
- Dan Lovinger
- Gianluigi Nusca
- Amos Ortal
- Ivan Pashov
- Ganesh Prasad
- Paul Swan
- Shiva Kumar Thangapandi

Amos Ortal and Dotan Elharrar were extremely helpful on NAP, and Shiva Kumar Thangapandi helped extensively with EAP.

The detailed checking Christophe Nasarre, overall technical reviewer, performed contributed greatly to the technical accuracy and consistency in the book.

We would like to again thank Ilfak Guilfanov of Hex-Rays (*www.hex-rays.com*) for the IDA Pro Advanced and Hex-Rays licenses they granted to Alex Ionescu so that he could speed up his reverse engineering of the Windows kernel.

Finally, the authors would like to thank the great staff at Microsoft Press who have been behind turning this book into a reality. Devon Musgrave served double duty as acquisitions editor and developmental editor, while Carol Dillingham oversaw the title as its project editor. Editorial and production manager Steve Sagman, copy editor Roger LeBlanc, proofreader Audrey Marr, and indexer Christina Yeager also contributed to the quality of this book.

Last but not least, thanks to Ben Ryan, publisher of Microsoft Press, who continues to believe in the importance of providing this level of detail about Windows to their readers!

Errata & Book Support

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

http://go.microsoft.com/FWLink/?Linkid=245675

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.*

Contents

Windows Internals, Sixth Edition, Part 1

Chapter 1 Concepts and Tools 1

Windows Operating System Versions 1
Foundation Concepts and Terms 2
Windows API 2
Services, Functions, and Routines 4
Processes, Threads, and Jobs 5
Virtual Memory 15
Kernel Mode vs. User Mode 17
Terminal Services and Multiple Sessions 20
Objects and Handles 21
Security 22
Registry 23
Unicode 24
Digging into Windows Internals 24
Performance Monitor 25
Kernel Debugging 26
Windows Software Development Kit 31
Windows Driver Kit 31
Sysinternals Tools 32
Conclusion 32

Chapter 2 System Architecture 33

Requirements and Design Goals 33
Operating System Model 34
Architecture Overview 35
Portability 37
Symmetric Multiprocessing 38

Scalability....40
Differences Between Client and Server Versions....41
Checked Build....45
Key System Components....46
Environment Subsystems and Subsystem DLLs....48
Ntdll.dll....53
Executive....54
Kernel....57
Hardware Abstraction Layer....60
Device Drivers....63
System Processes....68
Conclusion....78

Chapter 3 System Mechanisms 79

Trap Dispatching....79
Interrupt Dispatching....81
Timer Processing....112
Exception Dispatching....123
System Service Dispatching....132
Object Manager....140
Executive Objects....143
Object Structure....145
Synchronization....176
High-IRQL Synchronization....178
Low-IRQL Synchronization....183
System Worker Threads....205
Windows Global Flags....207
Advanced Local Procedure Call....209
Connection Model....210
Message Model....211
Asynchronous Operation....213
Views, Regions, and Sections....214
Attributes....215
Blobs, Handles, and Resources....215
Security....216
Performance....217
Debugging and Tracing....218

Kernel Event Tracing....220
Wow64....224
Wow64 Process Address Space Layout....224
System Calls....225
Exception Dispatching....225
User APC Dispatching....225
Console Support....225
User Callbacks....226
File System Redirection....226
Registry Redirection....227
I/O Control Requests....227
16-Bit Installer Applications....228
Printing....228
Restrictions....228
User-Mode Debugging....229
Kernel Support....229
Native Support....230
Windows Subsystem Support....232
Image Loader....232
Early Process Initialization....234
DLL Name Resolution and Redirection....235
Loaded Module Database....238
Import Parsing....242
Post-Import Process Initialization....243
SwitchBack....244
API Sets....245
Hypervisor (Hyper-V)....248
Partitions....249
Parent Partition....249
Child Partitions....251
Hardware Emulation and Support....254
Kernel Transaction Manager....268
Hotpatch Support....270
Kernel Patch Protection....272
Code Integrity....274
Conclusion....276

Chapter 4 Management Mechanisms **277**

The Registry 277
Viewing and Changing the Registry 277
Registry Usage 278
Registry Data Types 279
Registry Logical Structure 280
Transactional Registry (TxR) 287
Monitoring Registry Activity 289
Process Monitor Internals 289
Registry Internals 293
Services 305
Service Applications 305
The Service Control Manager 321
Service Startup 323
Startup Errors 327
Accepting the Boot and Last Known Good 328
Service Failures 330
Service Shutdown 331
Shared Service Processes 332
Service Tags 335
Unified Background Process Manager 336
Initialization 337
UBPM API 338
Provider Registration 338
Consumer Registration 339
Task Host 341
Service Control Programs 341
Windows Management Instrumentation 342
Providers 344
The Common Information Model and the Managed Object Format Language 345
Class Association 349
WMI Implementation 351
WMI Security 353
Windows Diagnostic Infrastructure 354
WDI Instrumentation 354
Diagnostic Policy Service 354
Diagnostic Functionality 356
Conclusion 357

Chapter 5 Processes, Threads, and Jobs 359

Process Internals 359
Data Structures 359
Protected Processes 368
Flow of *CreateProcess* 369
Stage 1: Converting and Validating Parameters and Flags 371
Stage 2: Opening the Image to Be Executed 373
Stage 3: Creating the Windows Executive Process Object (*PspAllocateProcess*) 376
Stage 4: Creating the Initial Thread and Its Stack and Context 381
Stage 5: Performing Windows Subsystem–Specific Post-Initialization 383
Stage 6: Starting Execution of the Initial Thread 385
Stage 7: Performing Process Initialization in the Context of the New Process 386
Thread Internals 391
Data Structures 391
Birth of a Thread 398
Examining Thread Activity 398
Limitations on Protected Process Threads 401
Worker Factories (Thread Pools) 403
Thread Scheduling 408
Overview of Windows Scheduling 408
Priority Levels 410
Thread States 416
Dispatcher Database 421
Quantum 422
Priority Boosts 430
Context Switching 448
Scheduling Scenarios 449
Idle Threads 453
Thread Selection 456
Multiprocessor Systems 458
Thread Selection on Multiprocessor Systems 467
Processor Selection 468
Processor Share-Based Scheduling 470
Distributed Fair Share Scheduling 471
CPU Rate Limits 478

Dynamic Processor Addition and Replacement 479
Job Objects 480
Job Limits 481
Job Sets 482
Conclusion 485

Chapter 6 Security **487**

Security Ratings 487
Trusted Computer System Evaluation Criteria 487
The Common Criteria 489
Security System Components 490
Protecting Objects 494
Access Checks 495
Security Identifiers 497
Virtual Service Accounts 518
Security Descriptors and Access Control 522
The AuthZ API 536
Account Rights and Privileges 538
Account Rights 540
Privileges 540
Super Privileges 546
Access Tokens of Processes and Threads 547
Security Auditing 548
Object Access Auditing 549
Global Audit Policy 552
Advanced Audit Policy Settings 554
Logon 555
Winlogon Initialization 556
User Logon Steps 558
Assured Authentication 562
Biometric Framework for User Authentication 563
User Account Control and Virtualization 566
File System and Registry Virtualization 566
Elevation 573
Application Identification (AppID) 581
AppLocker 583
Software Restriction Policies 589
Conclusion 590

Chapter 7 Networking **591**

Windows Networking Architecture 591
The OSI Reference Model 592
Windows Networking Components 594
Networking APIs 597
Windows Sockets 597
Winsock Kernel 603
Remote Procedure Call 605
Web Access APIs 610
Named Pipes and Mailslots 612
NetBIOS 618
Other Networking APIs 620
Multiple Redirector Support 627
Multiple Provider Router 627
Multiple UNC Provider 630
Surrogate Providers 632
Redirector 633
Mini-Redirectors 634
Server Message Block and Sub-Redirectors 635
Distributed File System Namespace 637
Distributed File System Replication 638
Offline Files 639
Caching Modes 641
Ghosts 643
Data Security 643
Cache Structure 643
BranchCache 645
Caching Modes 647
BranchCache Optimized Application Retrieval: SMB Sequence 651
BranchCache Optimized Application Retrieval: HTTP Sequence 653
Name Resolution 655
Domain Name System 655
Peer Name Resolution Protocol 656
Location and Topology 658
Network Location Awareness 658
Network Connectivity Status Indicator 659
Link-Layer Topology Discovery 662