TURING 图灵程序设计丛书 微软技术系列

Microsoft

Windows之父Allchin
Windows NT首席设计师Cutler 联袂推荐
微软公司副总裁Fathi

Windows Internals
Covering Windows Server 2008 and Windows Vista

5
FIFTH EDITION

深入解析Windows操作系统（第5版·英文版）

[美] Mark E. Russinovich
[美] David A. Solomon 著
[加] Alex Ionescu

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

人民邮电出版社
POSTS & TELECOM PRESS

TURING 图灵程序设计丛书 微软技术系列

Windows Internals

Covering Windows Server 2008 and Windows Vista Fifth Edition

深入解析Windows操作系统（第5版·英文版）

[美] Mark E. Russinovich
[美] David A. Solomon 著
[加] Alex Ionescu

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 = Windows Internals：Covering Windows Server 2008 and Windows Vista：第5版：英文 /（美）拉西诺维奇（Russinovich, M. E.），（美）所罗门（Solomon, D. A.），（加）艾欧内斯库（Ionescu, A.）著 . —北京：人民邮电出版社，2009.9
（图灵程序设计丛书）
ISBN 978-7-115-21165-1

Ⅰ. 深… Ⅱ. ①拉…②所…③艾… Ⅲ. 窗口软件，Windows－英文 Ⅳ. TP316.7

中国版本图书馆CIP数据核字（2009）第128683号

内 容 提 要

本书是操作系统内核专家 Mark Russinovich 和 David Solomon 的 Windows 操作系统原理的最新版著作，针对 Windows Vista 和 Windows Server 2008 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（第 5 版·英文版）

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

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
三河市海波印务有限公司印刷

◆ 开本：800×1000　1/16
印张：78.75
字数：1512千字　2009年 9 月第 1 版
印数：1－3 000册　2009年 9 月河北第 1 次印刷

著作权合同登记号　图字：01-2009-3816号
ISBN 978-7-115-21165-1/TP

定价：158.00元

读者服务热线：(010)51095186　印装质量热线：(010)67129223
反盗版热线：(010)67171154

站在巨人的肩上

Standing on Shoulders of Giants

www.turingbook.com

版 权 声 明

Foreword

It's both a pleasure and an honor for me to write the foreword for this latest edition of *Windows Internals*. Many significant changes have occurred in Windows since the last edition of the book, and David, Mark, and Alex have done an excellent job of updating the book to address them. Whether you are new to Windows internals or an old hand at kernel development, you will find lots of detailed analysis and examples to help improve your understanding of the core mechanisms of Windows as well as the general principles of operating system design.

Today, Windows enjoys unprecedented breadth and depth in the computing world. Variants of the original Windows NT design run on everything from Xbox game consoles to desktop and laptop computers to clusters of servers with dozens of processors and petabytes of storage. Advances such as hypervisors, 64-bit computing, multicore and many-core processor designs, flash-based storage, and wireless and peer-to-peer networking continue to provide plenty of interesting and innovative areas for operating system design.

One such area of innovation is security. Over the past decade, the entire computing industry—and Microsoft in particular—has been confronted with huge new threats, and security has become the top issue facing many of our customers. Attacks such as Blaster and Sasser threatened to bring the entire Internet to its knees, and Windows was at the eye of the hurricane. It was obvious to us that we could no longer afford to do business as usual, as many of the usability and simplicity features designed into Windows were being used to attack it for nefarious reasons. At first the hackers were teenagers trying to gain notoriety by breaking into systems or adding graffiti to a corporate Web site, but pretty soon the attacks intensified and went underground. The hackers became more sophisticated and evaded inspection. You rarely see headlines about viruses and worms these days, but make no mistake—botnets and identity theft are big business today, as are industrial and government espionage through targeted attacks.

In January 2002, Bill Gates sent his now-famous "Trustworthy Computing" memorandum to all Microsoft employees. It was a call to action that resonated well and charted the course for how we would build software and conduct business over the coming years. Nearly the entire Windows engineering team was diverted to work on Windows XP SP2, a service pack dedicated almost entirely to improving the security of the operating system. The Security Development Lifecycle (SDL) was developed and applied to all Microsoft products, with particular emphasis on Windows Vista as the first version of the operating system designed from the ground up to be secure. SDL specifies strict guidelines and processes for secure software development. Sophisticated tools have been developed to scan everything from source code to system binaries to network protocols for common security vulnerabilities. Every time a new security vulnerability is discovered, it is analyzed, and mitigations are developed to address that potential attack vector. Windows Vista has now been in the market for

two years, and it is by far the most secure version of Windows. Some industry analysts have pointed out that it is, in fact, the most secure general purpose operating system shipping today.

The Windows team has continued to innovate over the past few years. Windows XP, Windows Server 2003, Windows Server 2003 R2, Windows XP SP2, Windows Vista, Windows Server 2008, and Hyper-V are all major accomplishments and great successes—as well as great additions to the Windows family of products.

Frankly, I can't think of a more exciting and challenging topic. Nor can I think of a more authoritative and well-written book. David, Mark, and Alex have done a thorough job of dissecting the Windows architecture and providing diagnostic tools for hands-on learning. I hope you enjoy reading and learning about Windows as much as we all enjoy working on it.

Ben Fathi
Corporate Vice President, Windows Core Development
Microsoft Corporation

Introduction

Windows Internals, Fifth Edition is intended for advanced computer professionals (both developers and system administrators) who want to understand how the core components of the Windows Vista and Windows Server 2008 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

The first two chapters ("Concepts and Tools" and "System Architecture") lay the foundation with definitions and explanations of terms and concepts used throughout the rest of the book. The next two chapters—"System Mechanisms" and "Management Mechanisms"—describe key underlying mechanisms in the system. The next eight chapters explain the core components of the operating system: processes, threads, and jobs; security; the I/O system; storage management; memory management; the cache manager; file systems; and networking. The last two chapters cover startup and shutdown process and crash dump analysis.

History of the Book

This is the fifth 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.

Fifth Edition Changes

This latest edition has been updated to cover the kernel changes made in Windows Vista and Windows Server 2008. Hands-on experiments have been updated to reflect changes in tools, and newly added experiments use tools not available when the fourth edition was written. Additionally, content has been added to cover mechanisms that were not previously described, such as the image loader and user-mode debugging facility, and information about previously covered subjects has been expanded as well.

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 (*www.winsiderss.com*). 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 .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 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.

Find Additional Content Online

As new or updated material becomes available that complements this book, it will be posted online on the Microsoft Press Online Developer Tools Web site. The type of material you might find includes updates to book content, articles, links to companion content, errata, sample chapters, and more. This Web content is available at *www.microsoft.com/learning/books/online/developer* and is updated periodically.

Support

Every effort has been made to ensure the accuracy of this book. Should you run into any problems or issues, please refer to the sources listed below.

From the Authors

This book isn't perfect. No doubt it contains some inaccuracies, or possibly we've omitted some topics we should have covered. If you find anything you think is incorrect, or if you believe we should have included material that isn't here, please feel free to send e-mail to *winint@solsem.com*. Updates and corrections will be posted on the Web site *http://technet.microsoft.com/en-us/sysinternals/bb963901.aspx*.

From Microsoft Press

Microsoft Press provides corrections for books through the World Wide Web at the following address:

www.microsoft.com/mspress/support

Questions and Comments

In addition to sending feedback directly to the authors, if you have comments, questions, or ideas regarding the presentation or use of this book, you can send them to Microsoft using either of the following methods:

Postal mail:

Microsoft Press
Attn: Windows Internals Editor
One Microsoft Way
Redmond, WA 98052-6399

E-mail:

mspinput@microsoft.com

Please note that product support isn't offered through these mail addresses. For support information, visit Microsoft's Web site at *http://support.microsoft.com/*.

Acknowledgments

We dedicate this edition to **Jim Allchin**, our executive sponsor and champion before he retired from Microsoft. Jim supported our book work on this and earlier editions and was instrumental in bringing Mark Russinovich to Microsoft. In addition to shepherding Windows Vista out the door, Jim also oversaw the delivery of Windows 2000, Windows XP, and Windows Server 2003.

Each edition of this book has to acknowledge **Dave Cutler**, Senior Technical Fellow and the original architect of Windows NT. Dave originally approved David Solomon's source code access and has been supportive of his work to explain the internals of Windows through his training business as well as during the writing of the editions of this book.

We also thank three developers at Microsoft for contributing content that was incorporated into this edition:

- **Christian Allred,** who wrote detailed descriptions on transactional NTFS (TxF) internals, data structures, and behaviors
- **Stone Cong,** who wrote content and created diagrams about the Common Log File System (CLFS)
- **Adrian Marinescu,** who updated his heap manager section in the memory management chapter

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

Dmitry Anipko	Kwan Hyun	Ravi Mumulla	Jon Schwartz
Eugene Bak	Mehmet Iyigun	Adi Oltean	Valerie See
Karlito Bonnevie	Philippe Joubert	Vince Orgovan	Matt Setzer
Jon Cargille	Kwan Hyun Kim	Bernard Ourghanlian	Andrey Shedel
Dean DeWhitt	Kinshuman Kinshumann	Alexey Pakhunov	Neeraj Singh
Apurva Doshi	Alex Kirshenbaum	Milos Petrbok	Vikram Singh
Joseph East	Norbert Kusters	Daniel Pravat	Paul Sliwowicz
Tahsin Erdogan	Jeff Lambert	Ravi Pudipeddi	John Stephens
Cenk Ergan	Paul Leach	Melur Raghuraman	Deepu Thomas
Osman Ertugay	Scott Lee	Ramu Ramanathan	J. R. Tipton
Tom Fout	Mark Lloyd	Vlad Sadovsky	Davis Walker
Nar Ganapathy	Karan Mehra	Dragos Sambotin	Brad Waters
Robin Giese	Derek Moore	Jamie Schwartz	Bruce Worthington

Thanks also to Daniel Pearson (who teaches Windows internals for Dave Solomon) for his review and input.

Others might have contributed by answering questions in the hallway or cafeteria or by providing technical material—if we missed you, please forgive us!

The authors would like to thank Ilfak Guilfanov of Hex-Rays (*www.hex-rays.com*) for the IDA Pro Advanced and Hex-Rays licenses for Alex Ionescu for his use in speeding his reverse engineering of the Windows kernel. Alex chose not to have Windows source code access (as did Mark Russinovich before he joined Microsoft) to research the information for his work on this book, and these tools greatly facilitated his work. IDA's features turn reverse engineering into a powerful tool for understanding Windows internals. Combined with the Hex-Rays Decompiler, this analysis becomes even faster and more refined, as C code is directly presented instead of assembler, including all the right types.

Thanks also to Matt Ginzton of VMware, who arranged for Alex and David to receive VMware Workstation to use in their research for the book. VMware Workstation was used instead of Microsoft Virtual PC because of its support for 64-bit guests and multiple snapshots with nonpersisent disks. (These features are now supported by Hyper-V, Microsoft's new server virtualization offering, but at the time of writing, this support was not available).

Thanks to Mike Vance of AMD for providing Dave Solomon's AMD64 laptop for use in his book research and live classes.

Finally, we want to thank the team at Microsoft Press who helped turn this book from idea into reality:

- Ben Ryan (acquisitions editor at Microsoft Press) for shepherding another edition of this great book
- Kathleen Atkins (project editor) and Devon Musgrave (developmental editor) for launching and overseeing the project
- Andrea Fox (proofreader), Curtis Philips (project and production manager), and John Pierce (project editor and copyeditor) for laboriously going through all our chapters to tighten up text, find inconsistencies, and keep the manuscript to the high standards of Microsoft Press

Alex Ionescu, Mark Russinovich, and David Solomon
May 2009

Table of Contents