

# Contents

## Windows Internals, Sixth Edition, Part 1

---

(See appendix for Part 1's table of contents)

## Windows Internals, Sixth Edition, Part 2

---

<i>Introduction</i> .....	xv
<b>Chapter 8 I/O System</b> .....	<b>1</b>
I/O System Components .....	1
The I/O Manager .....	3
Typical I/O Processing .....	4
Device Drivers .....	5
Types of Device Drivers .....	5
Structure of a Driver .....	12
Driver Objects and Device Objects .....	14
Opening Devices .....	19
I/O Processing .....	25
Types of I/O .....	25
I/O Request to a Single-Layered Driver .....	33
I/O Requests to Layered Drivers .....	40
I/O Cancellation .....	48
I/O Completion Ports .....	53
I/O Prioritization .....	58
Container Notifications .....	65
Driver Verifier .....	65
Kernel-Mode Driver Framework (KMDF) .....	68
Structure and Operation of a KMDF Driver .....	68
KMDF Data Model .....	70
KMDF I/O Model .....	74

---

### 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](http://microsoft.com/learning/booksurvey)

User-Mode Driver Framework (UMDF) . . . . .	78
The Plug and Play (PnP) Manager . . . . .	81
Level of Plug and Play Support . . . . .	82
Driver Support for Plug and Play . . . . .	82
Driver Loading, Initialization, and Installation . . . . .	84
Driver Installation. . . . .	94
The Power Manager . . . . .	98
Power Manager Operation. . . . .	100
Driver Power Operation . . . . .	101
Driver and Application Control of Device Power . . . . .	105
Power Availability Requests . . . . .	105
Processor Power Management (PPM) . . . . .	108
Conclusion . . . . .	123

**Chapter 9 Storage Management 125**

Storage Terminology . . . . .	125
Disk Devices . . . . .	126
Rotating Magnetic Disks. . . . .	126
Solid State Disks . . . . .	128
Disk Drivers . . . . .	131
Winload . . . . .	132
Disk Class, Port, and Miniport Drivers . . . . .	132
Disk Device Objects . . . . .	136
Partition Manager . . . . .	138
Volume Management. . . . .	138
Basic Disks . . . . .	139
Dynamic Disks. . . . .	141
Multipartition Volume Management . . . . .	147
The Volume Namespace . . . . .	153
Volume I/O Operations. . . . .	159
Virtual Disk Service . . . . .	160
Virtual Hard Disk Support . . . . .	162
Attaching VHDs . . . . .	163
Nested File Systems . . . . .	163
BitLocker Drive Encryption . . . . .	163
Encryption Keys . . . . .	165
Trusted Platform Module (TPM) . . . . .	168
BitLocker Boot Process . . . . .	170
BitLocker Key Recovery . . . . .	172

Full-Volume Encryption Driver.....	173
BitLocker Management.....	174
BitLocker To Go.....	175
Volume Shadow Copy Service.....	177
Shadow Copies.....	177
VSS Architecture.....	177
VSS Operation.....	178
Uses in Windows.....	181
Conclusion.....	186

## **Chapter 10 Memory Management 187**

Introduction to the Memory Manager.....	187
Memory Manager Components.....	188
Internal Synchronization.....	189
Examining Memory Usage.....	190
Services Provided by the Memory Manager.....	193
Large and Small Pages.....	193
Reserving and Committing Pages.....	195
Commit Limit.....	199
Locking Memory.....	199
Allocation Granularity.....	199
Shared Memory and Mapped Files.....	200
Protecting Memory.....	203
No Execute Page Protection.....	204
Copy-on-Write.....	209
Address Windowing Extensions.....	210
Kernel-Mode Heaps (System Memory Pools).....	212
Pool Sizes.....	213
Monitoring Pool Usage.....	215
Look-Aside Lists.....	219
Heap Manager.....	220
Types of Heaps.....	221
Heap Manager Structure.....	222
Heap Synchronization.....	223
The Low Fragmentation Heap.....	223
Heap Security Features.....	224
Heap Debugging Features.....	225
Pageheap.....	226
Fault Tolerant Heap.....	227

Virtual Address Space Layouts . . . . .	228
x86 Address Space Layouts . . . . .	229
x86 System Address Space Layout . . . . .	232
x86 Session Space . . . . .	233
System Page Table Entries . . . . .	235
64-Bit Address Space Layouts . . . . .	237
x64 Virtual Addressing Limitations . . . . .	240
Dynamic System Virtual Address Space Management . . . . .	242
System Virtual Address Space Quotas . . . . .	245
User Address Space Layout . . . . .	246
Address Translation . . . . .	251
x86 Virtual Address Translation . . . . .	252
Translation Look-Aside Buffer . . . . .	259
Physical Address Extension (PAE) . . . . .	260
x64 Virtual Address Translation . . . . .	265
IA64 Virtual Address Translation . . . . .	266
Page Fault Handling . . . . .	267
Invalid PTEs . . . . .	268
Prototype PTEs . . . . .	269
In-Paging I/O . . . . .	271
Collided Page Faults . . . . .	272
Clustered Page Faults . . . . .	272
Page Files . . . . .	273
Commit Charge and the System Commit Limit . . . . .	275
Commit Charge and Page File Size . . . . .	278
Stacks . . . . .	279
User Stacks . . . . .	280
Kernel Stacks . . . . .	281
DPC Stack . . . . .	282
Virtual Address Descriptors . . . . .	282
Process VADs . . . . .	283
Rotate VADs . . . . .	284
NUMA . . . . .	285
Section Objects . . . . .	286
Driver Verifier . . . . .	292
Page Frame Number Database . . . . .	297
Page List Dynamics . . . . .	300
Page Priority . . . . .	310
Modified Page Writer . . . . .	314

PFN Data Structures . . . . .	315
Physical Memory Limits . . . . .	320
Windows Client Memory Limits . . . . .	321
Working Sets . . . . .	324
Demand Paging . . . . .	324
Logical Prefetcher . . . . .	324
Placement Policy . . . . .	328
Working Set Management . . . . .	329
Balance Set Manager and Swapper . . . . .	333
System Working Sets . . . . .	334
Memory Notification Events . . . . .	335
Proactive Memory Management (Superfetch) . . . . .	338
Components . . . . .	338
Tracing and Logging . . . . .	341
Scenarios . . . . .	342
Page Priority and Rebalancing . . . . .	342
Robust Performance . . . . .	344
ReadyBoost . . . . .	346
ReadyDrive . . . . .	348
Unified Caching . . . . .	348
Process Reflection . . . . .	351
Conclusion . . . . .	354

**Chapter 11 Cache Manager 355**

Key Features of the Cache Manager . . . . .	355
Single, Centralized System Cache . . . . .	356
The Memory Manager . . . . .	356
Cache Coherency . . . . .	356
Virtual Block Caching . . . . .	358
Stream-Based Caching . . . . .	358
Recoverable File System Support . . . . .	359
Cache Virtual Memory Management . . . . .	360
Cache Size . . . . .	361
Cache Virtual Size . . . . .	361
Cache Working Set Size . . . . .	361
Cache Physical Size . . . . .	363
Cache Data Structures . . . . .	364
Systemwide Cache Data Structures . . . . .	365
Per-File Cache Data Structures . . . . .	368

File System Interfaces . . . . .	373
Copying to and from the Cache . . . . .	374
Caching with the Mapping and Pinning Interfaces . . . . .	374
Caching with the Direct Memory Access Interfaces . . . . .	375
Fast I/O . . . . .	375
Read-Ahead and Write-Behind . . . . .	377
Intelligent Read-Ahead . . . . .	378
Write-Back Caching and Lazy Writing . . . . .	379
Write Throttling . . . . .	388
System Threads . . . . .	390
Conclusion . . . . .	390

## **Chapter 12 File Systems 391**

Windows File System Formats . . . . .	392
CDFS . . . . .	392
UDF . . . . .	393
FAT12, FAT16, and FAT32 . . . . .	393
exFAT . . . . .	396
NTFS . . . . .	397
File System Driver Architecture . . . . .	398
Local FSDs . . . . .	398
Remote FSDs . . . . .	400
File System Operation . . . . .	407
File System Filter Drivers . . . . .	413
Troubleshooting File System Problems . . . . .	415
Process Monitor Basic vs. Advanced Modes . . . . .	415
Process Monitor Troubleshooting Techniques . . . . .	416
Common Log File System . . . . .	416
NTFS Design Goals and Features . . . . .	424
High-End File System Requirements . . . . .	424
Advanced Features of NTFS . . . . .	426
NTFS File System Driver . . . . .	439
NTFS On-Disk Structure . . . . .	442
Volumes . . . . .	442
Clusters . . . . .	442
Master File Table . . . . .	443
File Record Numbers . . . . .	447
File Records . . . . .	447
File Names . . . . .	449

Resident and Nonresident Attributes . . . . .	453
Data Compression and Sparse Files . . . . .	456
The Change Journal File . . . . .	461
Indexing . . . . .	464
Object IDs . . . . .	466
Quota Tracking . . . . .	466
Consolidated Security . . . . .	467
Reparse Points . . . . .	469
Transaction Support . . . . .	469
NTFS Recovery Support . . . . .	477
Design . . . . .	478
Metadata Logging . . . . .	479
Recovery . . . . .	483
NTFS Bad-Cluster Recovery . . . . .	487
Self-Healing . . . . .	490
Encrypting File System Security . . . . .	491
Encrypting a File for the First Time . . . . .	494
The Decryption Process . . . . .	496
Backing Up Encrypted Files . . . . .	497
Copying Encrypted Files . . . . .	497
Conclusion . . . . .	498

**Chapter 13 Startup and Shutdown 499**

Boot Process . . . . .	499
BIOS Preboot . . . . .	499
The BIOS Boot Sector and Bootmgr . . . . .	502
The UEFI Boot Process . . . . .	512
Booting from iSCSI . . . . .	514
Initializing the Kernel and Executive Subsystems . . . . .	514
Smss, Csrss, and Wininit . . . . .	522
ReadyBoot . . . . .	527
Images That Start Automatically . . . . .	528
Troubleshooting Boot and Startup Problems . . . . .	529
Last Known Good . . . . .	530
Safe Mode . . . . .	530
Windows Recovery Environment (WinRE) . . . . .	534
Solving Common Boot Problems . . . . .	537
Shutdown . . . . .	542
Conclusion . . . . .	545

<b>Chapter 14 Crash Dump Analysis</b>	<b>547</b>
Why Does Windows Crash? . . . . .	547
The Blue Screen . . . . .	548
Causes of Windows Crashes . . . . .	549
Troubleshooting Crashes . . . . .	551
Crash Dump Files . . . . .	553
Crash Dump Generation . . . . .	559
Windows Error Reporting . . . . .	561
Online Crash Analysis . . . . .	563
Basic Crash Dump Analysis . . . . .	564
Notmyfault . . . . .	564
Basic Crash Dump Analysis . . . . .	565
Verbose Analysis . . . . .	567
Using Crash Troubleshooting Tools . . . . .	569
Buffer Overruns, Memory Corruption, and Special Pool . . . . .	569
Code Overwrite and System Code Write Protection . . . . .	573
Advanced Crash Dump Analysis . . . . .	574
Stack Trashes . . . . .	575
Hung or Unresponsive Systems . . . . .	577
When There Is No Crash Dump . . . . .	581
Analysis of Common Stop Codes . . . . .	585
0xD1 - DRIVER_IRQL_NOT_LESS_OR_EQUAL . . . . .	585
0x8E - KERNEL_MODE_EXCEPTION_NOT_HANDLED . . . . .	586
0x7F - UNEXPECTED_KERNEL_MODE_TRAP . . . . .	588
0xC5 - DRIVER_CORRUPTED_EXPOOL . . . . .	590
Hardware Malfunctions . . . . .	593
Conclusion . . . . .	594
<i>Appendix: Contents of Windows Internals, Sixth Edition, Part 1</i>	595
<i>Index</i>	603

---

**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](http://microsoft.com/learning/booksurvey)