Microsoft

Windows Internals

7
SEVENTH
EDITION

Part 1

System architecture, processes, threads, memory management, and more

Professional

Pavel Yosifovich
Alex Ionescu
Mark E. Russinovich
David A. Solomon

Microsoft

Windows Internals

Seventh Edition

Part 1

System architecture, processes, threads, memory management, and more

Pavel Yosifovich, Alex Ionescu,
Mark E. Russinovich, and David A. Solomon

Windows Internals, Part 1: System architecture, processes, threads, memory management, and more

Table of Contents

Cover

Title Page

Copyright Page

Acknowledgments

Contents

Introduction

Chapter 1 Concepts and tools

- Windows operating system versions
  - Windows 10 and future Windows versions
  - Windows 10 and OneCore
- Foundation concepts and terms
  - Windows API
  - Services, functions, and routines
  - Processes
  - Threads
  - Jobs
  - Virtual memory
  - Kernel mode vs. user mode
  - Hypervisor
  - Firmware
  - Terminal Services and multiple sessions

Table of Contents

Objects and handles

Security

Registry

Unicode

Digging into Windows internals

Performance Monitor and Resource Monitor

Kernel debugging

Windows Software Development Kit

Windows Driver Kit

Sysinternals tools

Conclusion

Chapter 2 System architecture

Requirements and design goals

Operating system model

Architecture overview

Portability

Symmetric multiprocessing

Scalability

Differences between client and server versions

Checked build

Virtualization-based security architecture overview

Key system components

Environment subsystems and subsystem DLLs

Other subsystems

Executive

Kernel

Hardware abstraction layer

Device drivers

Table of Contents

System processes

Conclusion

Chapter 3 Processes and jobs

Creating a process

CreateProcess* functions arguments

Creating Windows modern processes

Creating other kinds of processes

Process internals

Protected processes

Protected Process Light (PPL)

Third-party PPL support

Minimal and Pico processes

Minimal processes

Pico processes

Trustlets (secure processes)

Trustlet structure

Trustlet policy metadata

Trustlet attributes

System built-in Trustlets

Trustlet identity

Isolated user-mode services

Trustlet-accessible system calls

Flow of CreateProcess

Stage 1: Converting and validating parameters and flags

Stage 2: Opening the image to be executed

Stage 3: Creating the Windows executive process object

Stage 4: Creating the initial thread and its stack and context

Stage 5: Performing Windows subsystemspecific initialization

Table of Contents

Stage 6: Starting execution of the initial thread

Stage 7: Performing process initialization in the context of the new process

Terminating a process

Image loader

Early process initialization

DLL name resolution and redirection

Loaded module database

Import parsing

Post-import process initialization

SwitchBack

API Sets

Jobs

Job limits

Working with a job

Nested jobs

Windows containers (server silos)

Conclusion

Chapter 4 Threads

Creating threads

Thread internals

Data structures

Birth of a thread

Examining thread activity

Limitations on protected process threads

Thread scheduling

Overview of Windows scheduling

Priority levels

Thread states

Table of Contents

Dispatcher database

Quantum

Priority boosts

Context switching

Scheduling scenarios

Idle threads

Thread suspension

(Deep) freeze

Thread selection

Multiprocessor systems

Thread selection on multiprocessor systems

Processor selection

Heterogeneous scheduling (big.LITTLE)

Group-based scheduling

Dynamic fair share scheduling

CPU rate limits

Dynamic processor addition and replacement

Worker factories (thread pools)

Worker factory creation

Conclusion

Chapter 5 Memory management

Introduction to the memory manager

Memory manager components

Large and small pages

Examining memory usage

Internal synchronization

Services provided by the memory manager

Page states and memory allocations

Table of Contents

Commit charge and commit limit

Locking memory

Allocation granularity

Shared memory and mapped files

Protecting memory

Data Execution Prevention

Copy-on-write

Address Windowing Extensions

Kernel-mode heaps (system memory pools)

Pool sizes

Monitoring pool usage

Look-aside lists

Heap manager

Process heaps

Heap types

The NT heap

Heap synchronization

The low-fragmentation heap

The segment heap

Heap security features

Heap debugging features

Pageheap

Fault-tolerant heap

Virtual address space layouts

x86 address space layouts

x86 system address space layout

x86 session space

System page table entries

ARM address space layout

Table of Contents

- 64-bit address space layout
- x64 virtual addressing limitations
- Dynamic system virtual address space management
- System virtual address space quotas
- User address space layout

Address translation

- x86 virtual address translation
- Translation look-aside buffer
- x64 virtual address translation
- ARM virtual address translation

Page fault handling

- Invalid PTEs
- Prototype PTEs
- In-paging I/O
- Collided page faults
- Clustered page faults
- Page files
- Commit charge and the system commit limit
- Commit charge and page file size

Stacks

- User stacks
- Kernel stacks
- DPC stack

Virtual address descriptors

- Process VADs
- Rotate VADs

NUMA

Section objects

Table of Contents

Working sets

- Demand paging
- Logical prefetcher and ReadyBoot
- Placement policy
- Working set management
- Balance set manager and swapper
- System working sets
- Memory notification events

Page frame number database

- Page list dynamics
- Page priority
- Modified page writer and mapped page writer
- PFN data structures
- Page file reservation

Physical memory limits

- Windows client memory limits

Memory compression

- Compression illustration
- Compression architecture

Memory partitions

Memory combining

- The search phase
- The classification phase
- The page combining phase
- From private to shared PTE
- Combined pages release

Memory enclaves

- Programmatic interface

Table of Contents

Memory enclave initializations

Enclave construction

Loading data into an enclave

Initializing an enclave

Proactive memory management (SuperFetch)

Components

Tracing and logging

Scenarios

Page priority and rebalancing

Robust performance

ReadyBoost

ReadyDrive

Process reflection

Conclusion

Chapter 6 I/O system

I/O system components

The I/O manager

Typical I/O processing

Interrupt Request Levels and Deferred Procedure Calls

Interrupt Request Levels

Deferred Procedure Calls

Device drivers

Types of device drivers

Structure of a driver

Driver objects and device objects

Opening devices

I/O processing

Types of I/O

Table of Contents

- I/O request packets
- I/O request to a single-layered hardware-based driver
- I/O requests to layered drivers
- Thread-agnostic I/O
- I/O cancellation
- I/O completion ports
- I/O prioritization
- Container notifications

Driver Verifier

- I/O-related verification options
- Memory-related verification options

The Plug and Play manager

- Level of Plug and Play support
- Device enumeration
- Device stacks
- Driver support for Plug and Play
- Plug-and-play driver installation

General driver loading and installation

- Driver loading
- Driver installation

The Windows Driver Foundation

- Kernel-Mode Driver Framework
- User-Mode Driver Framework

The power manager

- Connected Standby and Modern Standby
- Power manager operation
- Driver power operation
- Driver and application control of device power
- Power management framework

Table of Contents

Power availability requests

Conclusion

Chapter 7 Security

Security ratings

Trusted Computer System Evaluation Criteria

The Common Criteria

Security system components

Virtualization-based security

Credential Guard

Device Guard

Protecting objects

Access checks

Security identifiers

Virtual service accounts

Security descriptors and access control

Dynamic Access Control

The AuthZ API

Conditional ACEs

Account rights and privileges

Account rights

Privileges

Super privileges

Access tokens of processes and threads

Security auditing

Object access auditing

Global audit policy

Advanced Audit Policy settings

AppContainers

Table of Contents

Overview of UWP apps

The AppContainer

Logon

Winlogon initialization

User logon steps

Assured authentication

Windows Biometric Framework

Windows Hello

User Account Control and virtualization

File system and registry virtualization

Elevation

Exploit mitigations

Process-mitigation policies

Control Flow Integrity

Security assertions

Application Identification

AppLocker

Software Restriction Policies

Kernel Patch Protection

PatchGuard

HyperGuard

Conclusion

Index