

# CRK (Curriculum Resource Kit)

Arkady Retik, DSc

Windows Academic Program Manager

Microsoft, Redmond

*Microsoft*

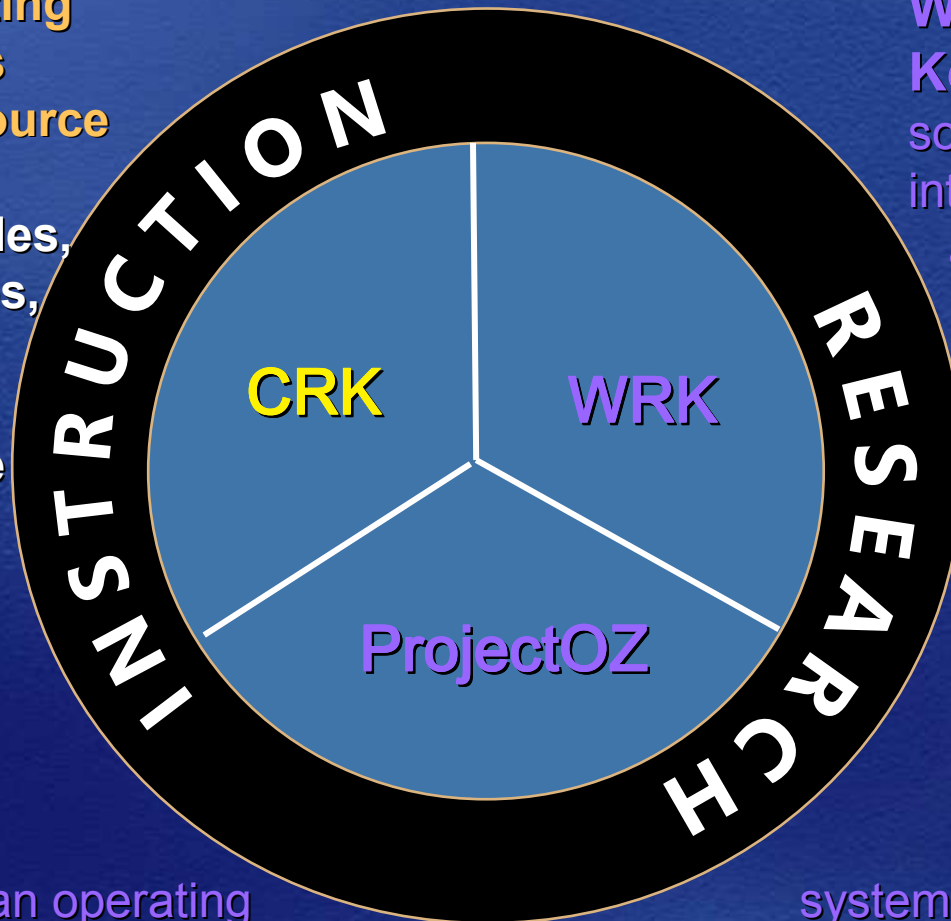
# Content

- Curriculum Resource Kit
  - Background
  - Overview
- Demo / hands-on
  - Units
  - Labs and Tools
  - Supplements
- Discussion and Q&As

# Windows Academic Program Components

## Windows Operating System Internals Curriculum Resource Kit (CRK) -

presentation slides, experiments, labs, quizzes and assignments for introducing case studies from the Windows kernel into operating system courses.



**Windows Research Kernel** – the core kernel sources and binaries integrated with an environment for building and testing experimental versions of the Windows kernel for use in teaching and research.

**ProjectOZ** - an operating systems project environment that uses the native kernel interfaces of Windows to provide simple, clean, user-mode abstractions of the CPU, MMU, trap mechanism, and physical memory that can be used to perform experiments in operating systems principles.



# CRK (Curriculum Resource Kit)

- Covers all ACM/IEEE/AIS OS BOK units and more (based on Windows XP/Server 2003)
- Scalable to multiple levels
- Modular (can be used in whole / in part)
- Case studies / compare & contrast

There are basic and advanced modules for each unit

❖ **Basic modules** provide materials to incorporate into a complete basic level OS course of one semester in length. The modules cover the Windows OS specific topics in the core and elective units of the OS BOK of Computing Curricula 2001.

❖ **Advanced modules** provide materials to incorporate into an advanced level OS course of one semester in length. The modules cover the Windows OS specific topics in the core and elective units of the "CC2001" OS BOK as well as in three supplementary units.

❑ contains references to source code

**Microsoft**

# CRK Units

## a. Core topics

- OS1. Overview of operating systems
- OS2. Operating system principles
- OS3. Concurrency
- OS4. Scheduling and dispatch
- OS5. Memory management

## b. Elective topics

- OS6. Device management
- OS7. Security and protection
- OS8. File systems
- OS9. Real-time and embedded systems
- OS10. Fault tolerance
- OS11. System perf evaluation & troubleshooting
- OS12. Scripting

## c. Supplementary topics

- OS-A. Windows networking
- OS-B. Comparing the Linux and Windows Kernels
- OS-C. Windows – Unix Interoperability

## d. Instructor Supplement - solutions to Quizzes, Programming Assignments

- Only available on your WAP CD and Faculty Connection portal !

### Most units contain

Labs & Exercises,  
Quizzes,  
Programming  
Assignments



# CRK Authors



## industry

**Mark Russinovich** is chief software architect and cofounder of Winternals Software ([www.winternals.com](http://www.winternals.com)), a company that specializes in advanced systems software for Microsoft Windows. Mark is co-author of *Inside Windows 2000*, 3rd Edition (Microsoft Press) with David Solomon and successor, *Windows Internals*, 4th Edition (Microsoft Press).

Mark is a Microsoft Most Valuable Professional (MVP) and serves as senior contributing editor for Windows IT Pro magazine where he contributes to the Windows Power Tools column. He is also a frequent speaker at major industry conferences such as Microsoft Tech Ed, IT Forum, Windows IT Pro Magazine's Connections and Redmond Magazine's TechMentor.

Mark has a B.S. from Carnegie Mellon University and a M.S. from Rensselaer Polytechnic Institute, both in computer engineering. In 1994, he earned a Ph.D. from Carnegie Mellon University, also in computer engineering.

**David Solomon** ([www.solsem.com](http://www.solsem.com)) teaches classes on Windows kernel internals to developers and IT professionals at companies worldwide, including Microsoft. He is the co-author of *Windows Internals*, 4th edition, the official Microsoft Press book on Windows kernel internals, as well as the previous edition, *Inside Windows 2000*. David also wrote *Inside Windows NT*, 2nd edition, and *Windows NT for OpenVMS Professionals*. He also co-created the Windows Internals COMPLETE video series which Microsoft licensed for worldwide internal training. David has served as technical chair for three past Windows NT conferences and has spoken at many TechEds and PDCs. He was a recipient of the 1993 & 2005 Microsoft Support Most Valuable Professional (MVP) award.



## academia



**Andreas Polze** is the Operating Systems and Middleware Professor at the Hasso-Plattner-Institute for Software Engineering at University Potsdam, Germany. He received a doctoral degree from Freie University Berlin, Germany, in 1994 and a habilitation degree from Humboldt University Berlin in 2001, both in computer science. His habilitation thesis investigates Predictable Computing in Multicomputer-Systems. Current research interests include Interconnecting Middleware and Embedded Systems, Mobility and Adaptive System Configuration, and End-to-End Service Availability for standard middleware platforms.

At University Potsdam, his current teaching activities focus on architecture of operating systems, on component-based middleware, as well as on predictable distributed computing. Our curriculum includes lectures that discuss operating system issues based on standard platforms (Windows 2000/XP, Mac OS X (BSD Unix), and Solaris) as well as on embedded systems (Windows CE, Embedded Linux).

Prof. Polze was a visiting scientist with the Dynamic Systems Unit at Software Engineering Institute, at Carnegie Mellon University, Pittsburgh, USA, where he worked on real-time computing on standard middleware (CORBA), and with the Real-Time Systems Laboratory at University of Illinois, Urbana-Champaign.

**Microsoft**

# CRK Contributors

- Many academics around the world participated in pilot review and provided comments and feedback:
  - US, Canada, Brazil, Mexico, Germany, UK, Russia, Israel, Australia, China, India
- CRK has been a top download since release in July 2006

**Thank you very much on behalf of Microsoft!!!**



# CRK Demo



Microsoft



# Demo & Hands-on Content

- CD overview
  - WAP – CRK, WRK, ProjectOZ
  - WinCE & PowerShell RC1
- CurriculumResourceKit–CRK folder
  - CRKUnits - 15 subfolders
    - Documents: Description, Syllabi, Lab-Setup, Glossary
  - CRKInstructorSupplement
  - CRKTools
  - WindowsInternalsBook4thEdition
- Q&As

# Demo & Hands-on Content – cont.

Each unit subfolder contains course material you can use and customize as you see fit:

- PowerPoint slides for:
  - lectures
  - additional optional demos of various tools
  - student labs
- student lab manual (Acrobat PDF file - created from the labs Powerpoint file)
- homework assignments (Microsoft Word DOC file)
- quizzes (Acrobat PDF file)

Many slides (especially labs) contain **notes**.

Solutions to homework assignments and quizzes are included with the CRK Instructor Supplement



# Demo & Hands-on Content – cont.

## Instructor materials:

The following icons are used on some slides:

- lab/exercise

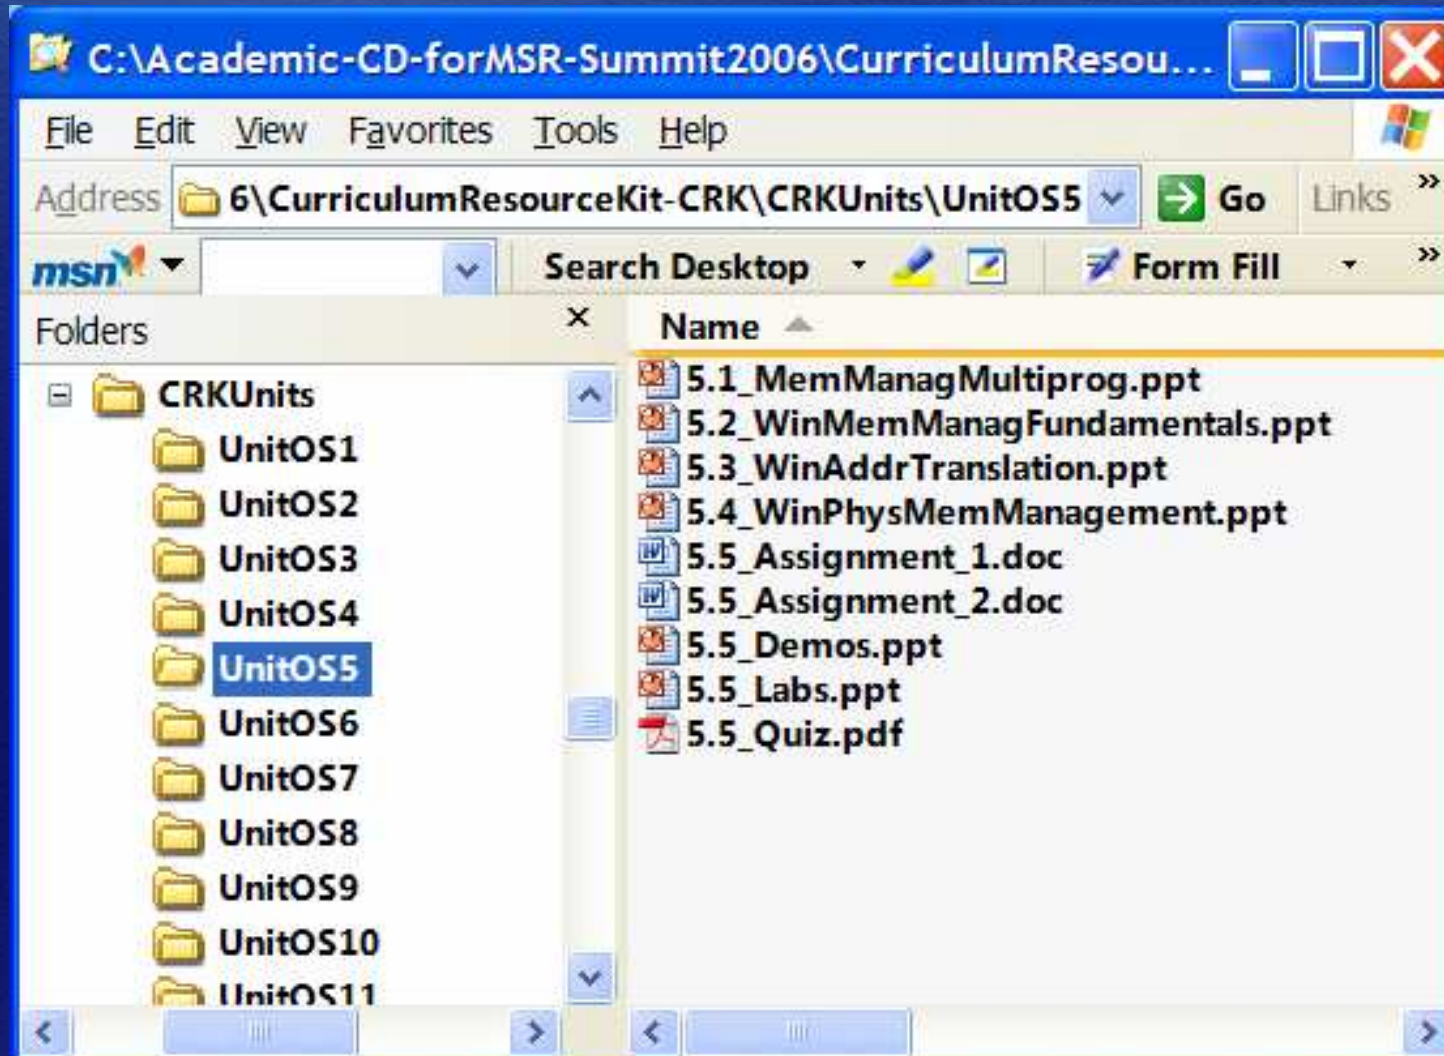


- informational slide



# Demo & Hands-on Content – cont.

- UnitOS5



Microsoft



# Unit OS5: Memory Management

## 5.4. Physical Memory Management

# Copyright Notice

© 2000-2005 David A. Solomon and Mark Russinovich

- These materials are part of the *Windows Operating System Internals Curriculum Development Kit*, developed by David A. Solomon and Mark E. Russinovich with Andreas Polze
- Microsoft has licensed these materials from David Solomon Expert Seminars, Inc. for distribution to academic organizations solely for use in academic environments (and not for commercial use)



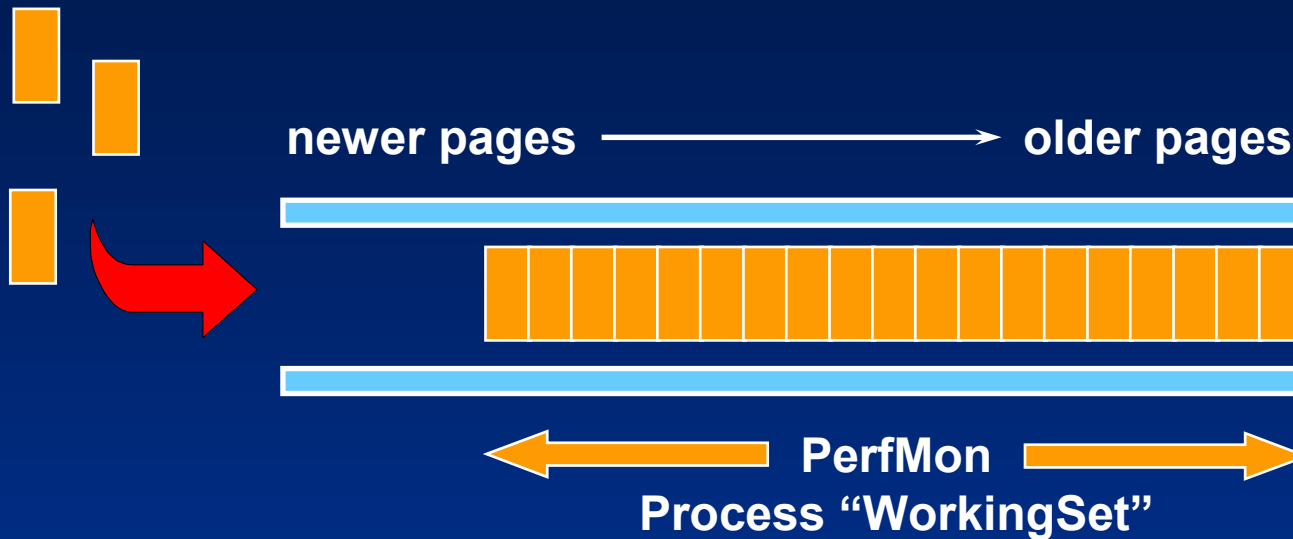
# Roadmap for Section 5.4.

- From working sets to paging dynamics
- Birth of a process working set
- Working set trimming, heuristics
- Paging, paging dynamics
- Hard vs. soft page faults
- Page files

# Working Set

- Working set: All the physical pages “owned” by a process
  - Essentially, all the pages the process can reference without incurring a page fault
- Working set limit: The maximum pages the process *can* own
  - When limit is reached, a page must be released for every page that’s brought in (“working set replacement”)
  - Default upper limit on size for each process
  - System-wide maximum calculated & stored in `MmMaximumWorkingSetSize`
    - approximately RAM minus 512 pages (2 MB on x86) minus min size of system working set (1.5 MB on x86)
    - Interesting to view (gives you an idea how much memory you’ve “lost” to the OS)
  - True upper limit: 2 GB minus 64 MB for 32-bit Windows

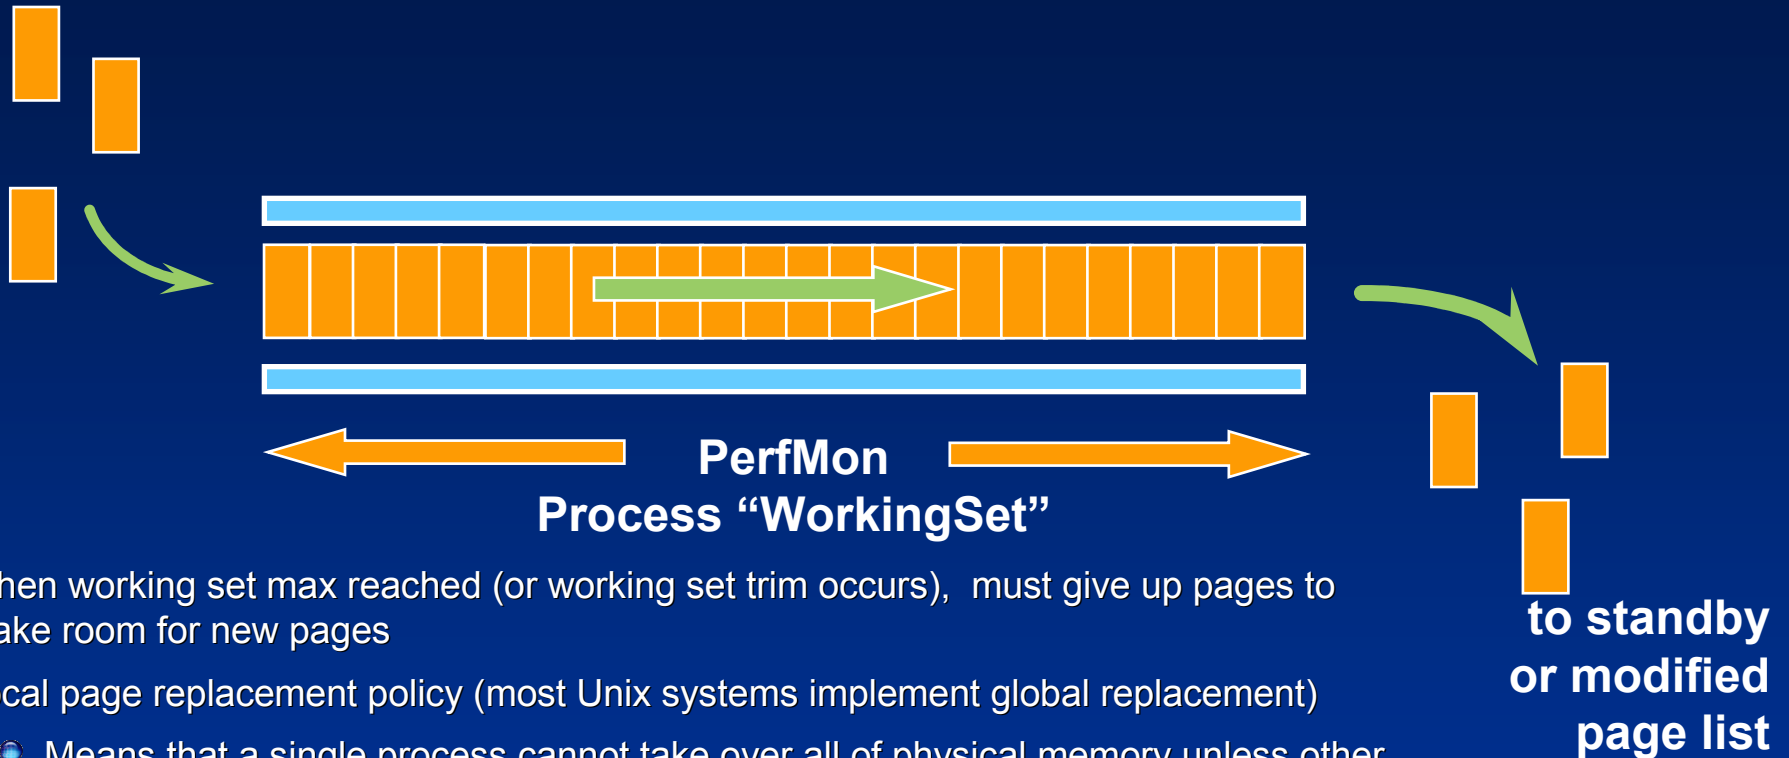
# Working Set List



- A process always starts with an empty working set
  - It then incurs *page faults* when referencing a page that isn't in its working set
  - Many page faults may be resolved from memory (to be described later)



# Working Set Replacement



- When working set max reached (or working set trim occurs), must give up pages to make room for new pages
- Local page replacement policy (most Unix systems implement global replacement)
  - Means that a single process cannot take over all of physical memory unless other processes aren't using it
- Page replacement algorithm is least recently accessed (pages are aged)
  - On UP systems only in Windows 2000 – done on all systems in Windows XP/Server 2003
- New VirtualAlloc flag in XP/Server 2003: MEM\_WRITE\_WATCH

# Soft vs. Hard Page Faults

- Types of “soft” page faults:
  - Pages can be faulted back into a process from the standby and modified page lists
  - A shared page that’s valid for one process can be faulted into other processes
- Some hard page faults unavoidable
  - Process startup (loading of EXE and DLLs)
  - Normal file I/O done via paging
    - Cached files are faulted into system working set
- To determine paging vs. normal file I/Os:
  - Monitor Memory->Page Reads/sec
    - Not Memory->Page Faults/sec, as that includes soft page faults
  - Subtract System->File Read Operations/sec from Page Reads/sec
  - Or, use Filemon to determine what file(s) are having paging I/O (asterisk next to I/O function)
  - Should not stay high for sustained period

# System Working Set

- Just as processes have working sets, Windows' pageable system-space code and data lives in the "system working set"
- Made up of 4 components:
  - Paged pool
  - Pageable code and data in the exec
  - Pageable code and data in kernel-mode drivers, Win32K.Sys, graphics drivers, etc.
  - Global file system data cache
- To get physical (resident) size of these with PerfMon, look at:
  - Memory | Pool Paged Resident Bytes
  - Memory | System Code Resident Bytes
  - Memory | System Driver Resident Bytes
  - **5** Memory | System Cache Resident Bytes
  - Memory | Cache bytes counter is total of these four "resident" (physical) counters (not just the cache; in NT4, same as "File Cache" on Task Manager / Performance tab)



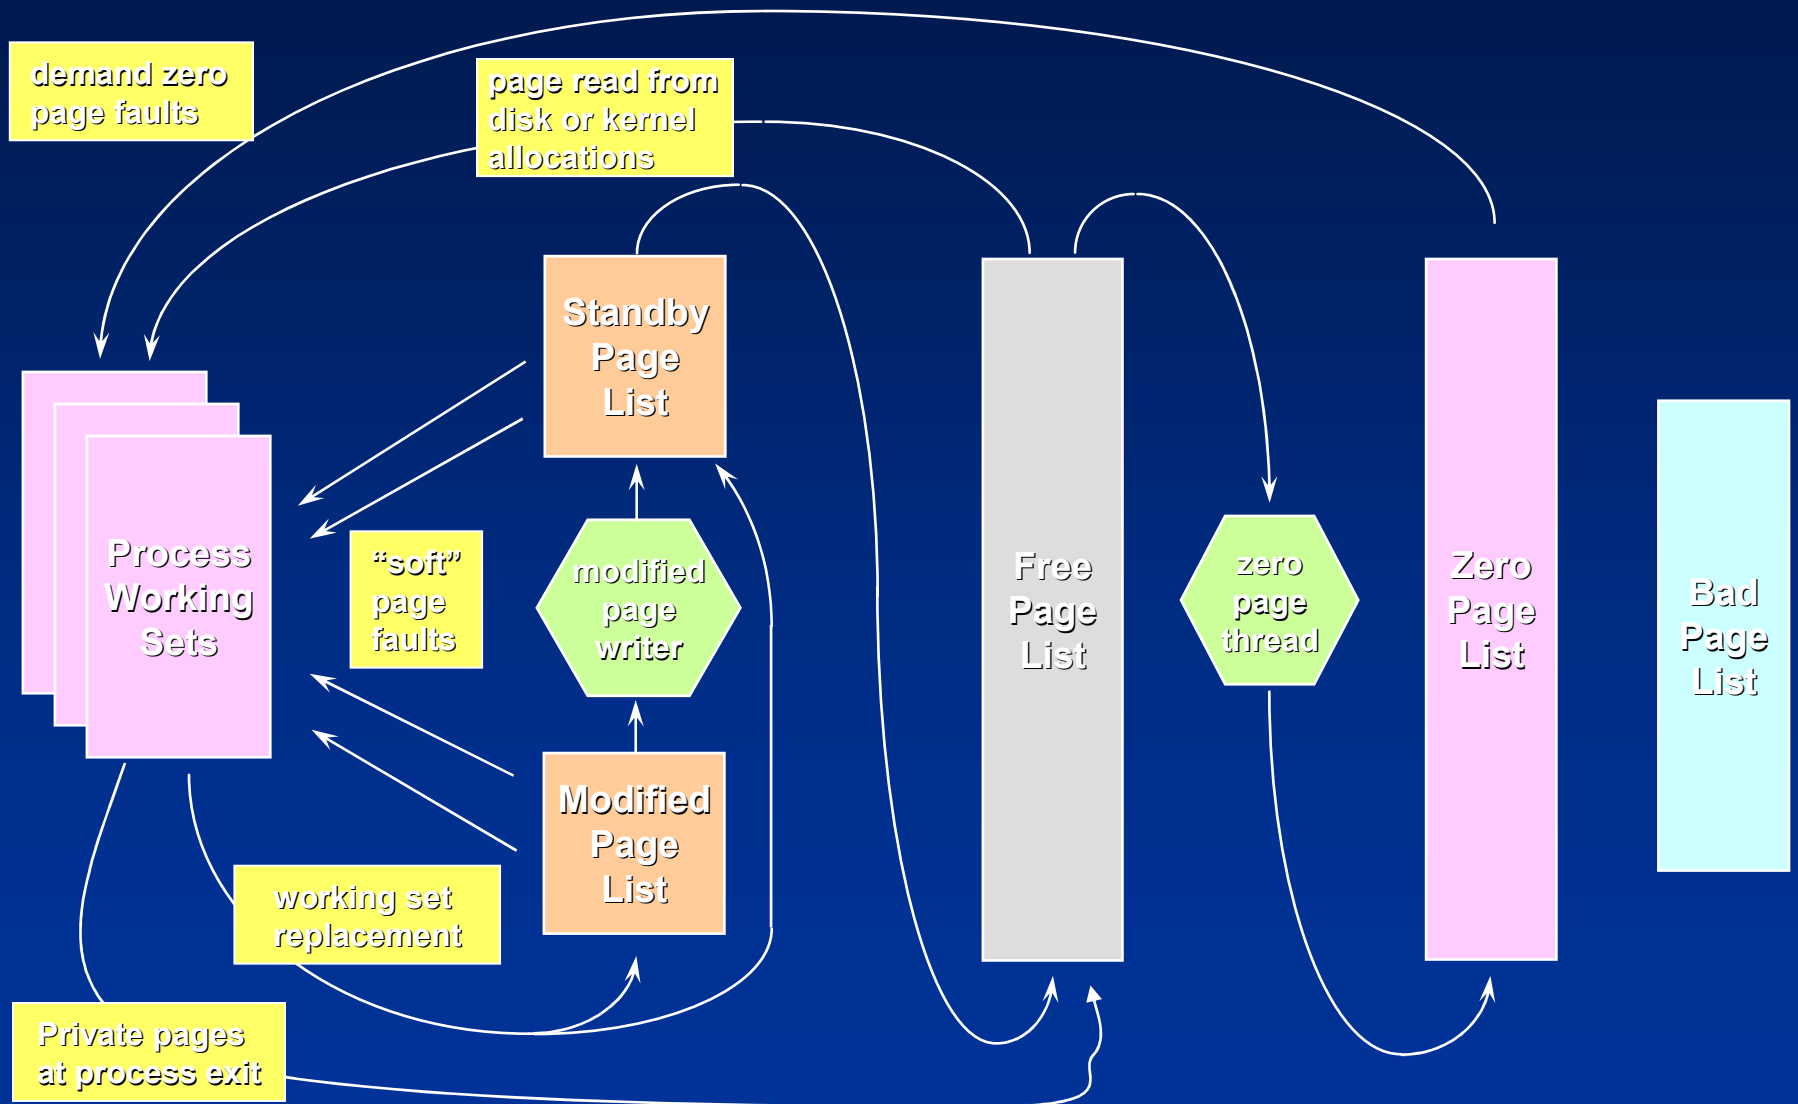
# Managing Physical Memory

- System keeps unassigned physical pages on one of several lists
  - Free page list
  - Modified page list
  - Standby page list
  - Zero page list
  - Bad page list - pages that failed memory test at system startup
- Lists are implemented by entries in the “PFN database”
  - Maintained as FIFO lists or queues

# Paging Dynamics


- New pages are allocated to working sets from the top of the free or zero page list
- Pages released from the working set due to working set replacement go to the bottom of:
  - The modified page list (if they were modified while in the working set)
  - The standby page list (if not modified)
    - Decision made based on “D” (dirty = modified) bit in page table entry
  - Association between the process and the physical page is still maintained while the page is on either of these lists

# Paging Dynamics






# Working Sets in Memory



Pages in Physical Memory



- As processes incur page faults, pages are removed from the free, modified, or standby lists and made part of the process working set
- A shared page may be resident in several processes' working sets at one time (this case not illustrated here)

# Page Frame Database – states of pages in physical memory

Status	Description
Active/valid	Page is part of working set (sys/proc), valid PTE points to it
Transition	Page not owned by a working set, not on any paging list I/O is in progress on this page
Standby	Page belonged to a working set but was removed; not modified
Modified	Removed from working set, modified, not yet written to disk
Modified no write	Modified page, will not be touched by modified page write, used by NTFS for pages containing log entries (explicit flushing)
Free	Page is free but has dirty data in it – cannot be given to user process – C2 security requirement
Zeroed	Page is free and has been initialized by zero page thread
Bad	Page has generated parity or other hardware errors

# Further Reading

- Mark E. Russinovich and David A. Solomon, Microsoft Windows Internals, 4th Edition, Microsoft Press, 2004.
- Chapter 7 - Memory Management
  - Page Fault Handling (from pp. 439)
  - Working Sets (from pp. 457)
  - Memory Pools (from pp. 399)
  - Page Frame Number Database (from pp. 469)



# Source Code References

- Windows Research Kernel sources
  - \base\ntos\mm – Memory manager
    - Wslist.c, Wsmanage.c – working set management
    - Pfnlist.c – physical memory list management
    - Modwrite.c – modified page writer
  - \base\ntos\inc\mm.h – additional structure definitions
  - \base\ntos\cache – Cache manager

# Demo & Hands-on Content: Tools

The tools are the ones referenced in the labs and demonstrations

SysInternals tools - about 50 tools and utilities

There are also 3 subfolders with these tools:

- \Debugging Tools - latest version of the Windows Debugging Tools

- \Reskit - old Windows 2000 Resource Kit tools not shipped in the 2003 Resource Kit

- \x64 - 64-bit versions of 3 tools for x64 (AMD64 and Intel EM64T - Itanium versions are not available)

- \Kernrate - Kernel Profiler tool referred to in Unit OS3

# Demo & Hands-on Content

## Sysinternals Tools

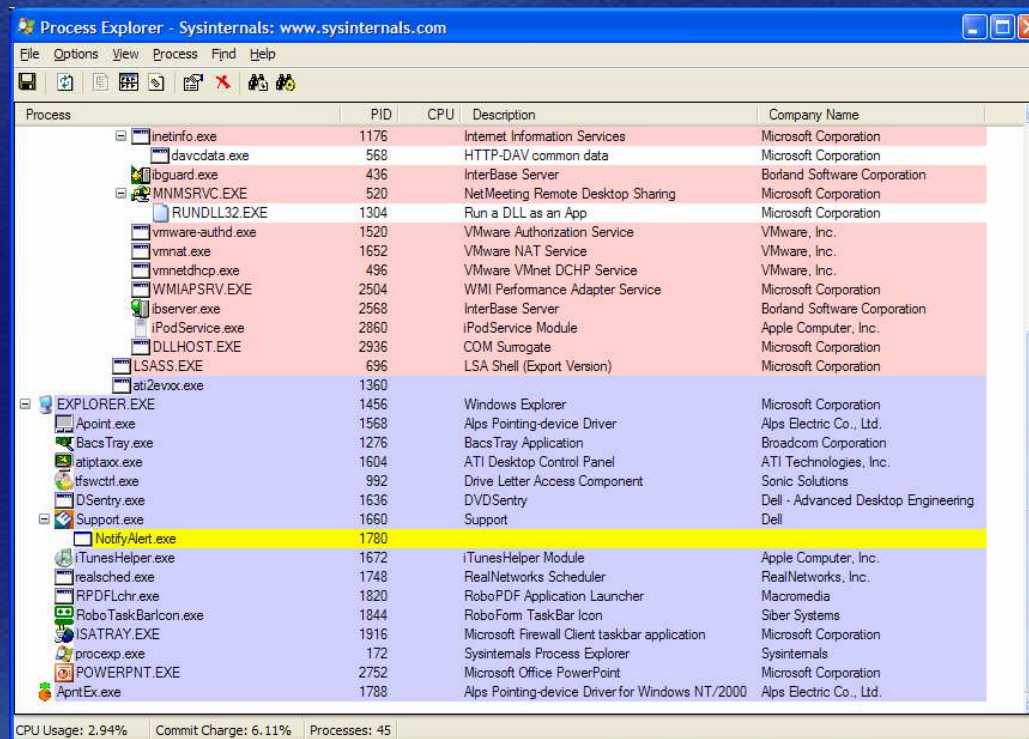
- System Information Tools
- Security Tools
- Development Tools
- Monitoring Tools
- Other Tools



# Demo & Hands-on Content: Sysinternals Tools

## Process Explorer

- “Super Task Manager”
- Runs on Win95, 98, ME, NT4, 2000, XP/2003, Vista



The screenshot shows the Process Explorer window with the following data:

Process	PID	CPU	Description	Company Name
inetinfo.exe	1176		Internet Information Services	Microsoft Corporation
davcddata.exe	568		HTTP-DAV common data	Microsoft Corporation
ibguard.exe	436		InterBase Server	Borland Software Corporation
WMNMSRVC.EXE	520		NetMeeting Remote Desktop Sharing	Microsoft Corporation
RUNDLL32.EXE	1304		Run a DLL as an App	Microsoft Corporation
vmware-authd.exe	1520		VMware Authorization Service	VMware, Inc.
vmnat.exe	1652		VMware NAT Service	VMware, Inc.
vmnetdhcp.exe	496		VMware VMnet DHCP Service	VMware, Inc.
WMIAPSRV.EXE	2504		WMI Performance Adapter Service	Microsoft Corporation
ibserver.exe	2568		InterBase Server	Borland Software Corporation
iPodService.exe	2860		iPodService Module	Apple Computer, Inc.
DLLHOST.EXE	2936		COM Surrogate	Microsoft Corporation
LSASS.EXE	696		LSA Shell (Export Version)	Microsoft Corporation
ati2evxx.exe	1360			
EXPLORER.EXE	1456		Windows Explorer	Microsoft Corporation
Apoint.exe	1568		Alps Pointing-device Driver	Alps Electric Co., Ltd.
Bacstray.exe	1276		Bacstray Application	Broadcom Corporation
atitptax.exe	1604		ATI Desktop Control Panel	ATI Technologies, Inc.
tfswctrl.exe	992		Drive Letter Access Component	Sonic Solutions
DSentry.exe	1636		DVD Sentry	Dell - Advanced Desktop Engineering
Support.exe	1660		Support	Dell
NotifyAlert.exe	1780			
iTunesHelper.exe	1672		iTunesHelper Module	Apple Computer, Inc.
realsched.exe	1748		RealNetworks Scheduler	RealNetworks, Inc.
RPDFLchr.exe	1820		RoboPDF Application Launcher	Macromedia
RoboTaskBarIcon.exe	1844		RoboForm TaskBar Icon	Siber Systems
ISATRAY.EXE	1916		Microsoft Firewall Client taskbar application	Microsoft Corporation
procepx.exe	172		Sysinternals Process Explorer	Sysinternals
POWERPNT.EXE	2752		Microsoft Office PowerPoint	Microsoft Corporation
ApntEx.exe	1788		Alps Pointing-device Driver for Windows NT/2000	Alps Electric Co., Ltd.

CPU Usage: 2.94% Commit Charge: 6.11% Processes: 45

Microsoft



# Demo & Hands-on Content:

## Sysinternals Tools

### PS Tools

- Psfile – lists & closes remote file opens
- Psshutdown – remote shutdown, lock workstation, log off user
- Psexec – run an app on a remote system
- Pslist – list processes & threads
- Psuptime – system up time
- Psinfo – display general system info
- Psgetsid – displays computer or user SIDs
- Psservice – service process control (like SC in XP)
- Psloglist – dumps event log in text
- PsSuspend – suspend a process
- PsKill – kill processes
- Psloggedon – lists local and remote logon sessions
- Pspassword – change local/remote passwords

# Curriculum Resource Kit

- Q&A
- Discussion:
  - More content?
  - More labs?
  - Video?
  - Academic Textbook?
  - Local language?

# Curriculum Resource Kit: Summary

- **Units and Tools** are available for free download from Academic Alliance Repository on MSDN  
[www.msdnaacr.net/curriculum/pfv.aspx?ID=6191](http://www.msdnaacr.net/curriculum/pfv.aspx?ID=6191)
- **Instructor Supplement** for faculty only, password protected access
  - Available for download to MSDNAA departmental subscribers  
<http://msdn.microsoft.com/subscriptions/>
  - Individual faculty can get access via Faculty Connection portal  
<http://www.microsoft.com/education/facultyconnection/>
- **Feedback** - Contact us at [compsci@microsoft.com](mailto:compsci@microsoft.com)
- Curriculum **Forum** for discussion and Q&A available on  
<http://forums.microsoft.com/WindowsAcademic>