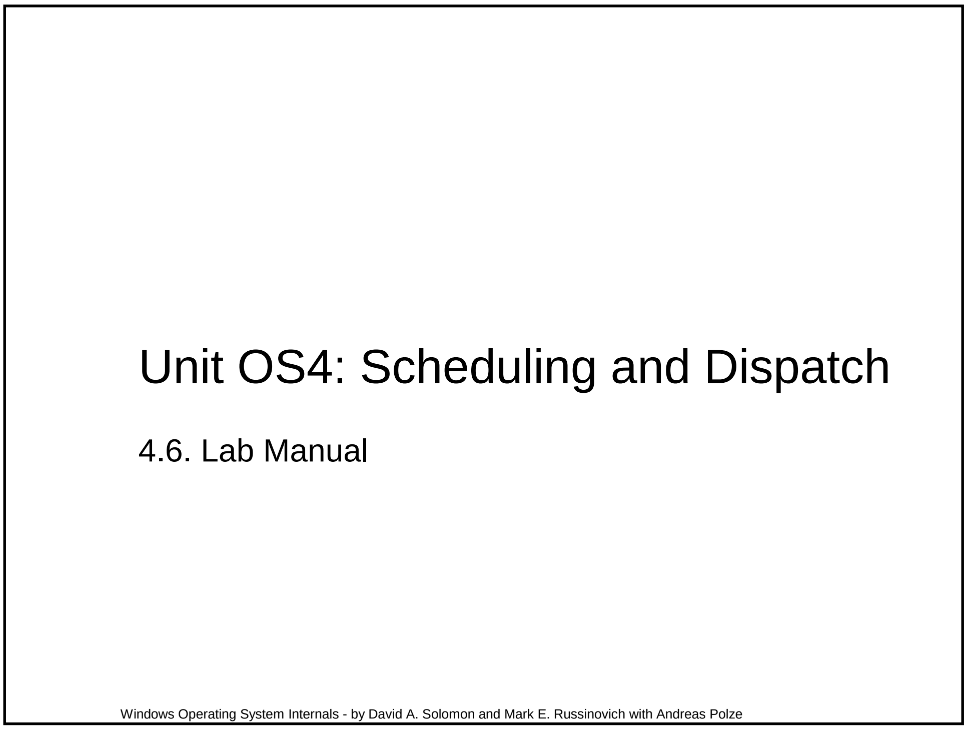

Unit OS4: Scheduling and Dispatch

4.6. Lab Manual

Windows Operating System Internals - by David A. Solomon and Mark E. Russinovich with Andreas Polze

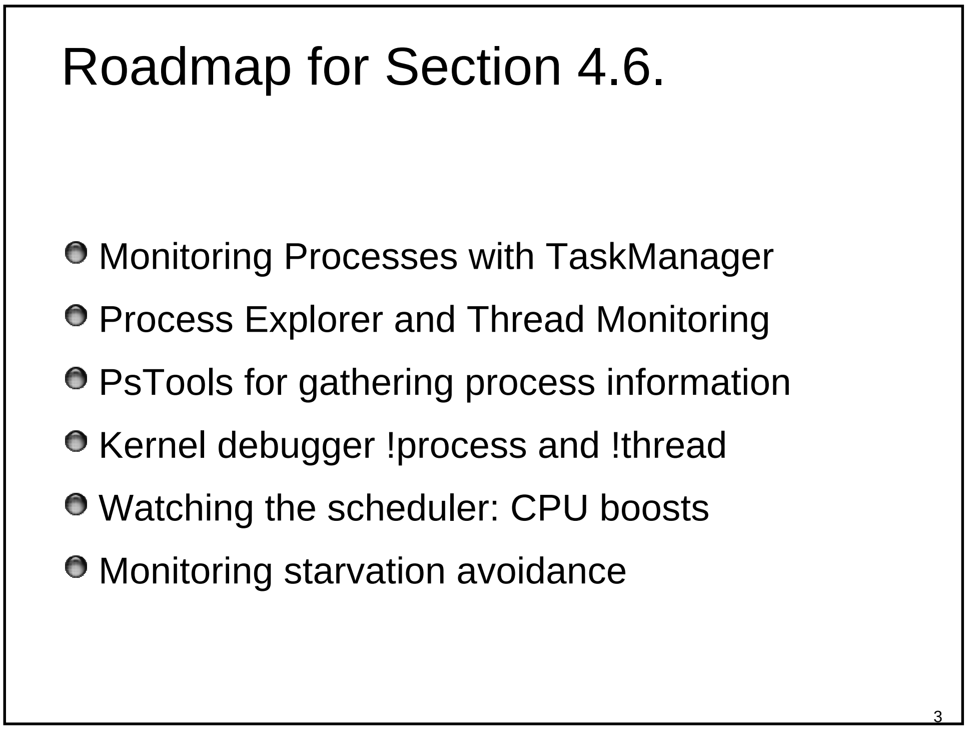

Roadmap for Section 4.6.

- Monitoring Processes with TaskManager
- Process Explorer and Thread Monitoring
- PsTools for gathering process information
- Kernel debugger !process and !thread
- Watching the scheduler: CPU boosts
- Monitoring starvation avoidance

3

Task Manager: Processes vs Applications Tabs

- Processes tab: List of processes
- Applications tab: List of top level visible windows

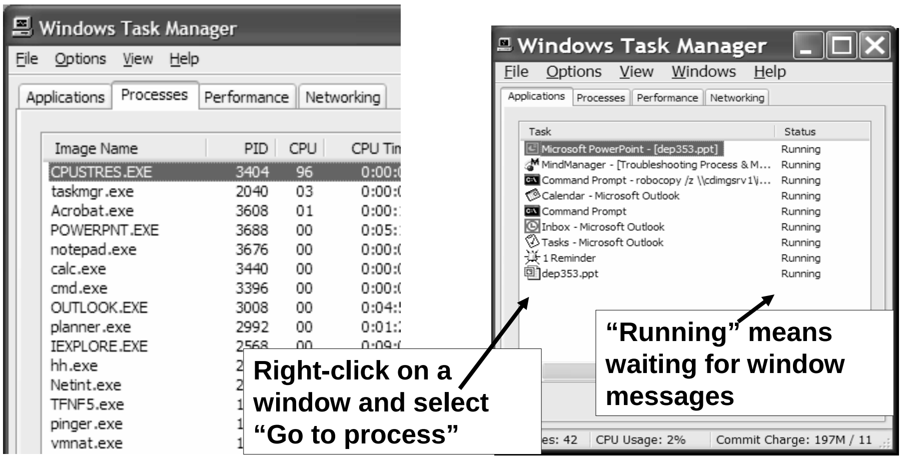
Windows Task Manager
File Options View Help
Applications Processes Performance Networking

Image Name	PID	CPU	CPU Tin
CPUSTRES.EXE	3404	96	0:00:
taskmgr.exe	2040	03	0:00:
Acrobat.exe	3608	01	0:00:
POWERPNT.EXE	3688	00	0:05:
notepad.exe	3676	00	0:00:
calc.exe	3440	00	0:00:
cmd.exe	3396	00	0:00:
OUTLOOK.EXE	3008	00	0:04:
planner.exe	2992	00	0:01:
IEXPLORE.EXE	2568	00	0:09:
hh.exe	2		
Netint.exe	2		
TFNF5.exe	1		
pinger.exe	1		
vmnat.exe	1		

Windows Task Manager
File Options View Windows Help
Applications Processes Performance Networking

Task	Status
Microsoft PowerPoint - [dep353.ppt]	Running
MindManager - [Troubleshooting Process & M...	Running
Command Prompt - robocopy /z \\cdimgsrv1\j...	Running
Calendar - Microsoft Outlook	Running
Command Prompt	Running
Inbox - Microsoft Outlook	Running
Tasks - Microsoft Outlook	Running
1 Reminder	Running
dep353.ppt	Running

Right-click on a window and select "Go to process"

"Running" means waiting for window messages

Understand Task Managers "Applications"

- A meaningless term at the OS level
  - Not a list of processes
  - Not a list of "tasks" (another meaningless term)
  - It's a list of top level visible windows in your session that meet certain criteria
- What does the status column mean?
  - Running:
    - Windows don't run—threads do
    - Running displayed only when owning thread is waiting for a window message (e.g. not running!)
  - Not Responding: not waiting for window messages
- To map a window to a process, right-click on a window and select "Go to process"

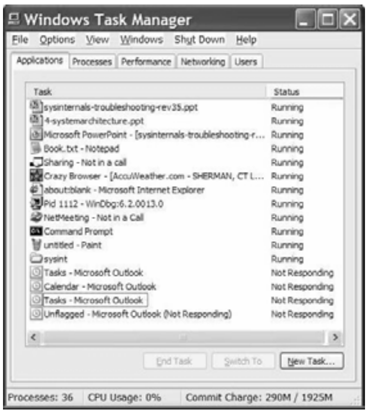
Windows Task Manager
File Options View Windows Shut Down Help
Applications Processes Performance Networking Users

Task	Status
sysinternals-troubleshooting-rev35.ppt	Running
4-systemarchitecture.ppt	Running
Microsoft PowerPoint - [sysinternals-troubleshooting-r...	Running
Book.txt - Notepad	Running
Sharing - Not in a call	Running
Crazy Browser - [AccuWeather.com - SHERMAN, CT L...	Running
about:blank - Microsoft Internet Explorer	Running
Pid 1112 - WinDbg:6.2.0013.0	Running
NetMeeting - Not in a Call	Running
Command Prompt	Running
untitled - Paint	Running
sysint	Running
Tasks - Microsoft Outlook	Not Responding
Calendar - Microsoft Outlook	Not Responding
Tasks - Microsoft Outlook	Not Responding
Unflagged - Microsoft Outlook (Not Responding)	Not Responding

Process Explorer (Sysinternals)

- "Super Task Manager"
  - Shows full image path, command line, environment variables, parent process, security access token, open handles, loaded DLLs & mapped files

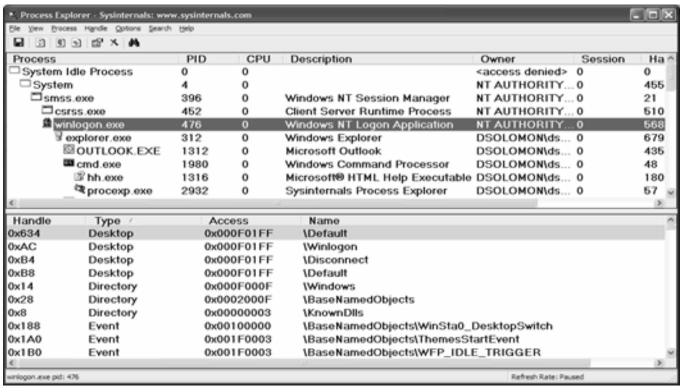
Process Explorer - Sysinternals: www.sysinternals.com
File View Process Handle Options Search Help

Process	PID	CPU	Description	Owner	Session	Ha
System Idle Process	0	0		<access denied>	0	0
System	4	0		NT AUTHORITY...	0	455
smss.exe	396	0	Windows NT Session Manager	NT AUTHORITY...	0	21
csrss.exe	452	0	Client Server Runtime Process	NT AUTHORITY...	0	510
winlogon.exe	476	0	Windows NT Logon Application	NT AUTHORITY...	0	568
explorer.exe	312	0	Windows Explorer	DSOLOMON\ds...	0	679
OUTLOOK.EXE	1312	0	Microsoft Outlook	DSOLOMON\ds...	0	435
cmd.exe	1980	0	Windows Command Processor	DSOLOMON\ds...	0	48
hh.exe	1316	0	Microsoft® HTML Help Executable	DSOLOMON\ds...	0	180
procexp.exe	2932	0	Sysinternals Process Explorer	DSOLOMON\ds...	0	57

Handle	Type	Access	Name
0x634	Desktop	0x000F01FF	\Default
0xAC	Desktop	0x000F01FF	\Winlogon
0xB4	Desktop	0x000F01FF	\Disconnect
0xB8	Desktop	0x000F01FF	\Default
0x14	Directory	0x000F000F	\Windows
0x28	Directory	0x0002000F	\BaseNamedObjects
0x8	Directory	0x00000003	\KnownDlls
0x188	Event	0x00100000	\BaseNamedObjects\WinSta0_DesktopSwitch
0x1A0	Event	0x001F0003	\BaseNamedObjects\ThemesStartEvent
0x1B0	Event	0x001F0003	\BaseNamedObjects\WFP_IDLE_TRIGGER

winlogon.exe pid: 476

Refresh Rate: Paused

Process Explorer's Process List

1. Run Process Explorer & maximize window
2. Run Task Manager – click on Processes tab
3. Arrange windows so you can see both
4. Notice process tree vs flat list in Task Manager
   - If parent has exited, process is left justified
5. Sort on first column ("Process") and note tree view disappears
6. Click on View->Show Process Tree (or CTRL+T) to bring it back
7. Notice description and company name columns
8. Hover mouse over image to see full path of image
9. Right click on a process and choose "Google"

Process Performance

- Click on Performance Tab of process properties
  - Note: all these numbers can be configured as columns

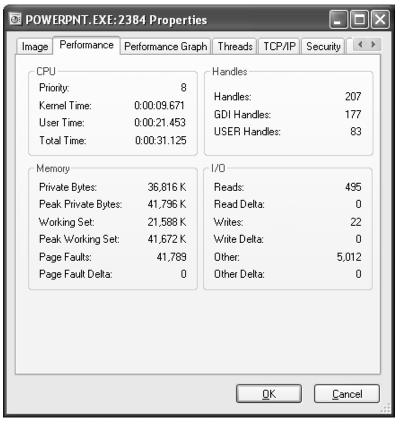

CPU

Priority: 8
Kernel Time: 0:00:09.671
User Time: 0:00:21.453
Total Time: 0:00:31.125

Handles

Handles: 207
GDI Handles: 177
USER Handles: 83

Memory

Private Bytes: 36,816 K
Peak Private Bytes: 41,796 K
Working Set: 21,588 K
Peak Working Set: 41,672 K
Page Faults: 41,789
Page Fault Delta: 0

I/O

Reads: 495
Read Delta: 0
Writes: 22
Write Delta: 0
Other: 5,012
Other Delta: 0

8

Thread Details

- Process Explorer "Threads" tab shows which thread(s) are running
  - Start address represents where the thread began running (not where it is now)
  - Click Module to get details on module containing thread start address

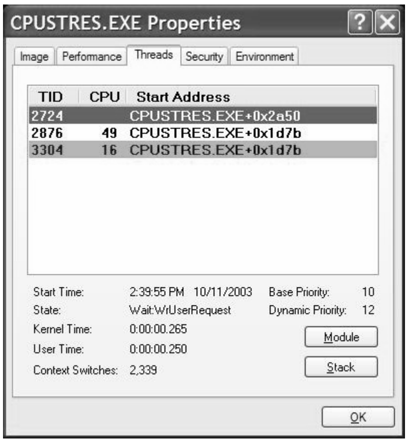

TID	CPU	Start Address
2724		CPUSTRES.EXE+0x2a50
2876	49	CPUSTRES.EXE+0x1d7b
3304	16	CPUSTRES.EXE+0x1d7b

Start Time: 2:39:55 PM 10/11/2003
State: Wait:WrUserRequest
Kernel Time: 0:00:00.265
User Time: 0:00:00.250
Context Switches: 2,339
Base Priority: 10
Dynamic Priority: 12

9

Thread Start Functions

- Process Explorer can map the addresses within a module to the names of functions
  - This can help identify which component within a process is responsible for CPU usage
- Requires access to:
  - Symbol file for that module
  - Proper version of Dbghelp.dll
- By default, Process Explorer looks for:
  - Dbghelp.dll: in the default Windows Debugging Tools install directory
  - Symbols: _NT_SYMBOL_PATH environment variable
  - Can also specify with Options->Configure Symbols

Process Explorer Lab: Environment Variables

- Click on Environment Tab of process properties

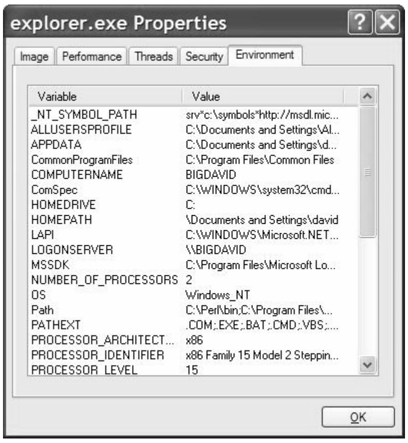

Variable	Value
_NT_SYMBOL_PATH	srv*c:\symbols*http://msdl.mic...
ALLUSERSPROFILE	C:\Documents and Settings\Al...
APPDATA	C:\Documents and Settings\d...
CommonProgramFiles	C:\Program Files\Common Files
COMPUTERNAME	BIGDAVID
ComSpec	C:\WINDOWS\system32\cmd...
HOMEDRIVE	C:
HOMEPATH	\Documents and Settings\david
LAPI	C:\WINDOWS\Microsoft.NET...
LOGONSERVER	\\BIGDAVID
MSSDK	C:\Program Files\Microsoft Lo...
NUMBER_OF_PROCESSORS	2
OS	Windows_NT
Path	C:\Perl\bin;C:\Program Files\...
PATHEXT	.COM;.EXE;.BAT;.CMD;.VBS;...
PROCESSOR_ARCHITECT...	x86
PROCESSOR_IDENTIFIER	x86 Family 15 Model 2 Steppin...
PROCESSOR_LEVEL	15

Identify Jobs used by WMI

- Jobs are used by WMI
  - Example: run Psinfo (Sysinternals) and pause output

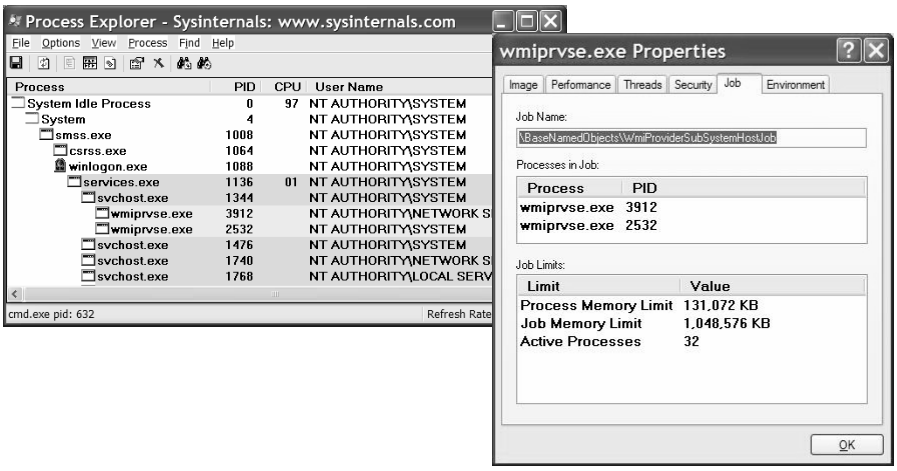

Process Explorer - Sysinternals: www.sysinternals.com
File Options View Process Find Help

Process	PID	CPU	User Name
System Idle Process	0	97	NT AUTHORITY\SYSTEM
System	4		NT AUTHORITY\SYSTEM
smss.exe	1008		NT AUTHORITY\SYSTEM
csrss.exe	1064		NT AUTHORITY\SYSTEM
winlogon.exe	1088		NT AUTHORITY\SYSTEM
services.exe	1136	01	NT AUTHORITY\SYSTEM
svchost.exe	1344		NT AUTHORITY\SYSTEM
wmiprvse.exe	3912		NT AUTHORITY\NETWORK S
wmiprvse.exe	2532		NT AUTHORITY\SYSTEM
svchost.exe	1476		NT AUTHORITY\SYSTEM
svchost.exe	1740		NT AUTHORITY\NETWORK S
svchost.exe	1768		NT AUTHORITY\LOCAL SERV

cmd.exe pid: 632

Job Name: \BaseNamedObjects\WmiProviderSubSystemHostJob

Processes in Job:

Process	PID
wmiprvse.exe	3912
wmiprvse.exe	2532

Job Limits:

Limit	Value
Process Memory Limit	131,072 KB
Job Memory Limit	1,048,576 KB
Active Processes	32

Jobs created by RUNAS

1. In a command prompt:
   RUNAS /USER:xxx CMD
   (where xxx is some other local account)
2. In ProcExp, find newly created cmd.exe process
   - Who is the father?
3. Run Notepad from new CMD window
4. Double click on newly highlighted process & click on Job tab

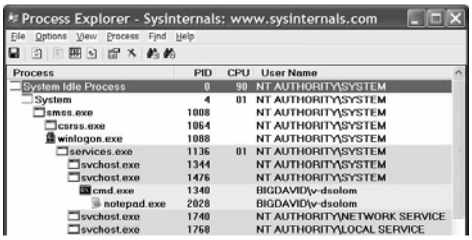

Process Explorer - Sysinternals: www.sysinternals.com
File Options View Process Find Help

Process	PID	CPU	User Name
System Idle Process	0	90	NT AUTHORITY\SYSTEM
System	4	01	NT AUTHORITY\SYSTEM
smss.exe	1008		NT AUTHORITY\SYSTEM
csrss.exe	1064		NT AUTHORITY\SYSTEM
winlogon.exe	1088		NT AUTHORITY\SYSTEM
services.exe	1136	01	NT AUTHORITY\SYSTEM
svchost.exe	1344		NT AUTHORITY\SYSTEM
svchost.exe	1476		NT AUTHORITY\SYSTEM
cmd.exe	1340		BIGDAVID\v-dsolom
notepad.exe	2028		BIGDAVID\v-dsolom
svchost.exe	1740		NT AUTHORITY\NETWORK SERVICE
svchost.exe	1768		NT AUTHORITY\LOCAL SERVICE

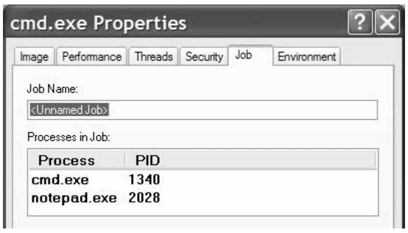

Job Name: <Unnamed Job>

Processes in Job:

Process	PID
cmd.exe	1340
notepad.exe	2028

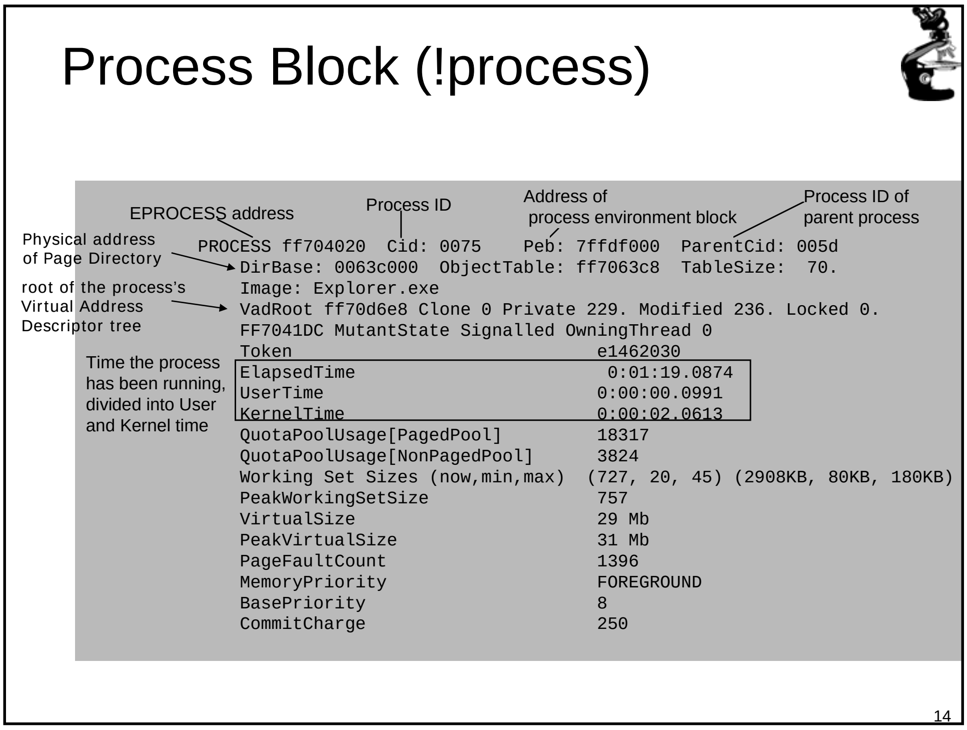

Process Block (!process)

EPROCESS address

Process ID

Address of process environment block

Process ID of parent process

Physical address of Page Directory

root of the process's Virtual Address Descriptor tree

Time the process has been running, divided into User and Kernel time

```
PROCESS ff704020  Cid: 0075    Peb: 7ffdf000  ParentCid: 005d
    DirBase: 0063c000  ObjectTable: ff7063c8  TableSize:  70.
    Image: Explorer.exe
    VadRoot ff70d6e8 Clone 0 Private 229. Modified 236. Locked 0.
    FF7041DC MutantState Signalled OwningThread 0
    Token                             e1462030
    ElapsedTime                        0:01:19.0874
    UserTime                          0:00:00.0991
    KernelTime                        0:00:02.0613
    QuotaPoolUsage[PagedPool]         18317
    QuotaPoolUsage[NonPagedPool]      3824
    Working Set Sizes (now,min,max)  (727, 20, 45) (2908KB, 80KB, 180KB)
    PeakWorkingSetSize                757
    VirtualSize                       29 Mb
    PeakVirtualSize                   31 Mb
    PageFaultCount                    1396
    MemoryPriority                    FOREGROUND
    BasePriority                      8
    CommitCharge                      250
```

14

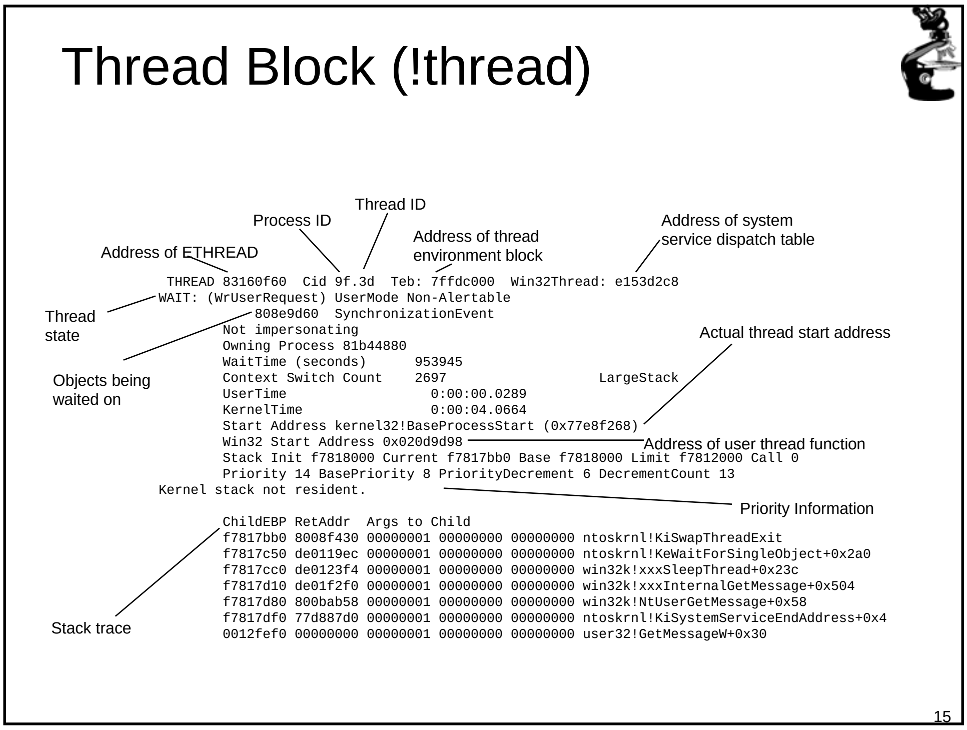

Thread Block (!thread)

Thread ID

Process ID

Address of ETHREAD

Address of thread environment block

Address of system service dispatch table

Thread state

Objects being waited on

Actual thread start address

Address of user thread function

Priority Information

Stack trace

```
    THREAD 83160f60  Cid 9f.3d  Teb: 7ffdc000  Win32Thread: e153d2c8
   WAIT: (WrUserRequest) UserMode Non-Alertable
               808e9d60  SynchronizationEvent
           Not impersonating
           Owning Process 81b44880
           WaitTime (seconds)      953945
           Context Switch Count    2697                   LargeStack
           UserTime                  0:00:00.0289
           KernelTime                0:00:04.0664
           Start Address kernel32!BaseProcessStart (0x77e8f268)
           Win32 Start Address 0x020d9d98
           Stack Init f7818000 Current f7817bb0 Base f7818000 Limit f7812000 Call 0
           Priority 14 BasePriority 8 PriorityDecrement 6 DecrementCount 13
   Kernel stack not resident.

           ChildEBP RetAddr  Args to Child
           f7817bb0 8008f430 00000001 00000000 00000000 ntoskrnl!KiSwapThreadExit
           f7817c50 de0119ec 00000001 00000000 00000000 ntoskrnl!KeWaitForSingleObject+0x2a0
           f7817cc0 de0123f4 00000001 00000000 00000000 win32k!xxxSleepThread+0x23c
           f7817d10 de01f2f0 00000001 00000000 00000000 win32k!xxxInternalGetMessage+0x504
           f7817d80 800bab58 00000001 00000000 00000000 win32k!NtUserGetMessage+0x58
           f7817df0 77d887d0 00000001 00000000 00000000 ntoskrnl!KiSystemServiceEndAddress+0x4
           0012fef0 00000000 00000001 00000000 00000000 user32!GetMessageW+0x30
```

15

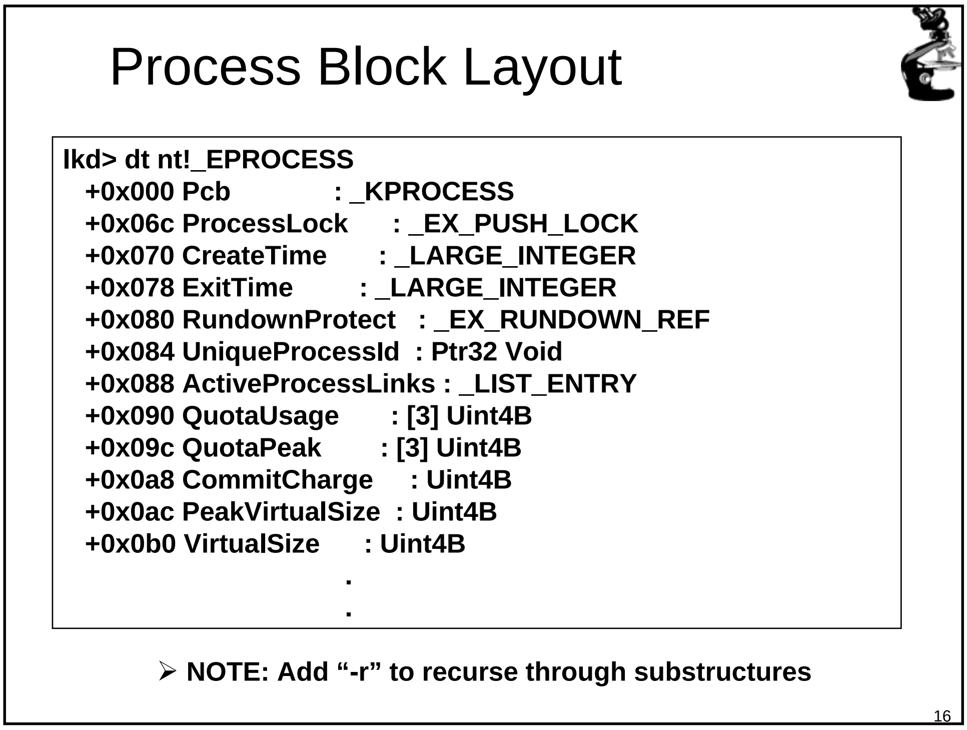

Process Block Layout

```
lkd> dt nt!_EPROCESS
   +0x000 Pcb              : _KPROCESS
   +0x06c ProcessLock      : _EX_PUSH_LOCK
   +0x070 CreateTime       : _LARGE_INTEGER
   +0x078 ExitTime         : _LARGE_INTEGER
   +0x080 RundownProtect   : _EX_RUNDOWN_REF
   +0x084 UniqueProcessId  : Ptr32 Void
   +0x088 ActiveProcessLinks : _LIST_ENTRY
   +0x090 QuotaUsage       : [3] Uint4B
   +0x09c QuotaPeak        : [3] Uint4B
   +0x0a8 CommitCharge     : Uint4B
   +0x0ac PeakVirtualSize  : Uint4B
   +0x0b0 VirtualSize      : Uint4B
                         .
                         .
```

➢ **NOTE: Add “-r” to recurse through substructures**

16

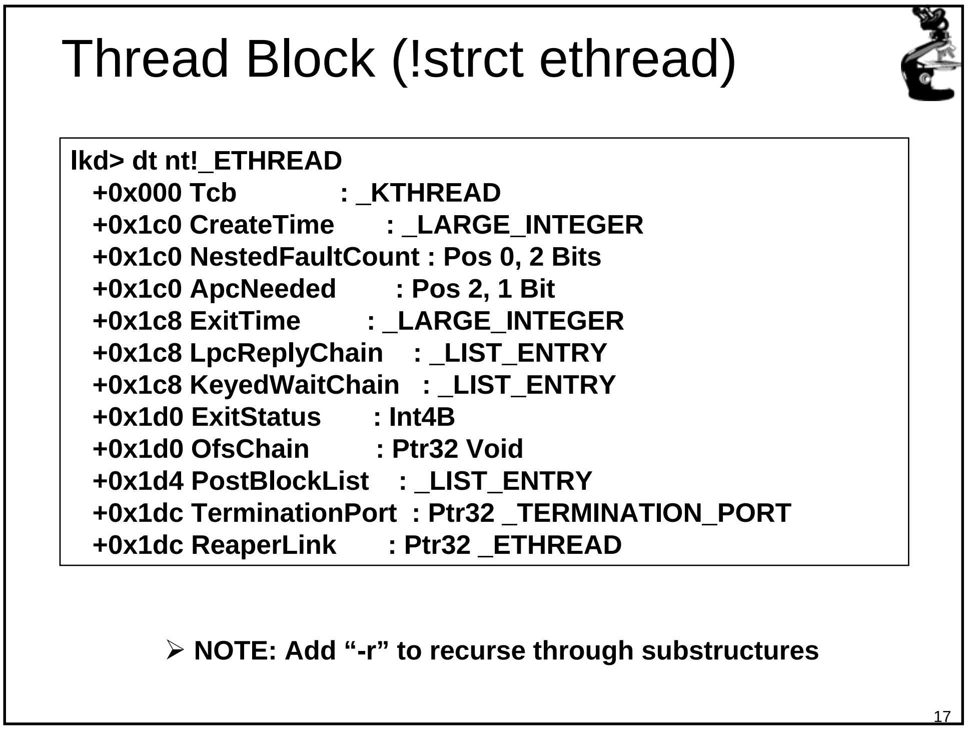

Thread Block (!strct ethread)

```
lkd> dt nt!_ETHREAD
   +0x000 Tcb              : _KTHREAD
   +0x1c0 CreateTime       : _LARGE_INTEGER
   +0x1c0 NestedFaultCount : Pos 0, 2 Bits
   +0x1c0 ApcNeeded        : Pos 2, 1 Bit
   +0x1c8 ExitTime         : _LARGE_INTEGER
   +0x1c8 LpcReplyChain    : _LIST_ENTRY
   +0x1c8 KeyedWaitChain   : _LIST_ENTRY
   +0x1d0 ExitStatus       : Int4B
   +0x1d0 OfsChain         : Ptr32 Void
   +0x1d4 PostBlockList    : _LIST_ENTRY
   +0x1dc TerminationPort  : Ptr32 _TERMINATION_PORT
   +0x1dc ReaperLink       : Ptr32 _ETHREAD
```

➢ **NOTE: Add “-r” to recurse through substructures**

17

Watching the Scheduler

Performance Monitor - Options | Chart

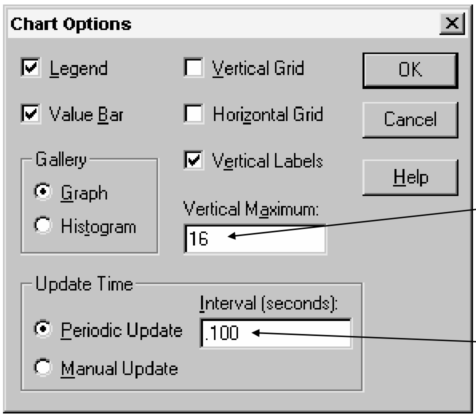

Chart Options

☑ Legend

☑ Value Bar

Gallery
● Graph
○ Histogram

Update Time
● Periodic Update
○ Manual Update

Interval (seconds): .100

□ Vertical Grid

□ Horizontal Grid

☑ Vertical Labels

Vertical Maximum: 16

Set chart maximum vertical scale to 16

Set update interval to 0.1 seconds or less

Screen snapshot from: Performance Monitor Options menu | Chart command

Watching the Scheduler (contd.)

Performance Monitor

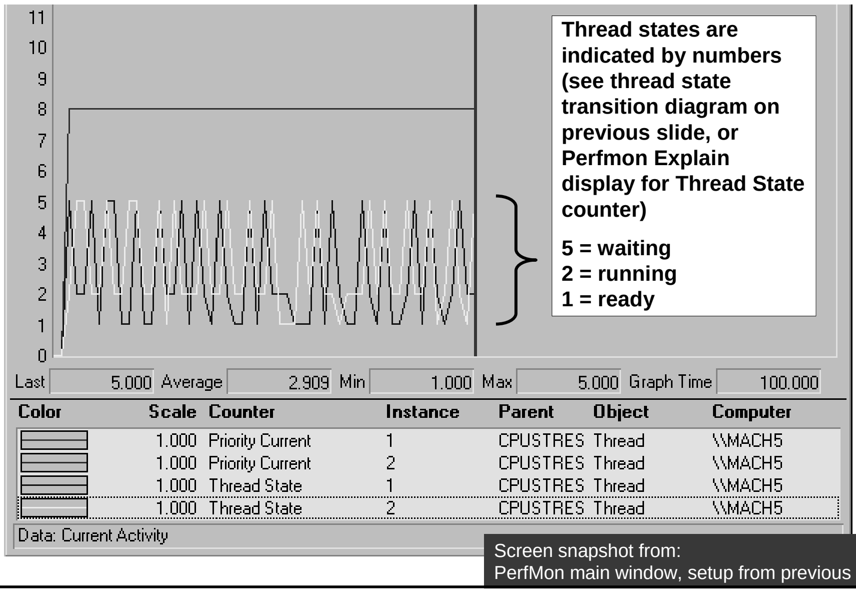

Last 5.000 Average 2.909 Min 1.000 Max 5.000 Graph Time 100.000

Color	Scale	Counter	Instance	Parent	Object	Computer
▭	1.000	Priority Current	1	CPUSTRES	Thread	\\MACH5
▭	1.000	Priority Current	2	CPUSTRES	Thread	\\MACH5
▭	1.000	Thread State	1	CPUSTRES	Thread	\\MACH5
▭	1.000	Thread State	2	CPUSTRES	Thread	\\MACH5

Data: Current Activity

Thread states are indicated by numbers (see thread state transition diagram on previous slide, or Perfmon Explain display for Thread State counter)

5 = waiting
2 = running
1 = ready

Screen snapshot from:
PerfMon main window, setup from previous slide

Watching Forground Priority Boosts

- Run: cpustres.exe (Resource Kit)

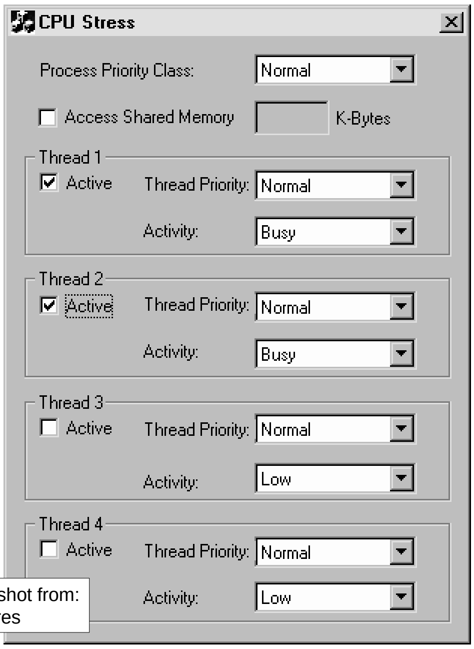

CPU Stress

Process Priority Class: Normal

☐ Access Shared Memory K-Bytes

Thread 1
☑ Active Thread Priority: Normal
Activity: Busy

Thread 2
☑ Active Thread Priority: Normal
Activity: Busy

Thread 3
☐ Active Thread Priority: Normal
Activity: Low

Thread 4
☐ Active Thread Priority: Normal
Activity: Low

Screen snapshot from:
Run… cpustres

Priority Boost and Decay (contd.)

Demo with CpuStres and PerfMon

- CpuStres settings:
  - two active threads
  - activity level = busy (about 25% wait time)
  - normal process priority class, normal thread priorities
- Usually only visible in PerfMon if target app owns foreground window (hence longer quantum)
- These are showing +2 boost (from 8 to 10) for foreground apps after wait completion

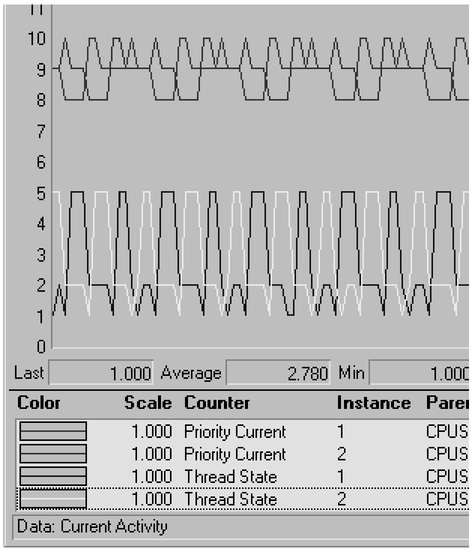

Last 1.000 Average 2.780 Min 1.000

Color	Scale	Counter	Instance	Parei
▭	1.000	Priority Current	1	CPUS
▭	1.000	Priority Current	2	CPUS
▭	1.000	Thread State	1	CPUS
▭	1.000	Thread State	2	CPUS

Data: Current Activity

Priority Boosts on GUI Threads

- Threads that own windows receive an additional boost of 2 when they wake up because of windowing activity, such as the arrival of window messages.
- The windowing system (Win32k.sys) applies this boost when it calls KeSetEvent to set an event used to wake up a GUI thread.
- The reason for this boost is similar to the previous one—to favor interactive applications.

CPU Starvation Resolution

- CpuStres with two compute-bound threads (“maximum” activity level)
- One is at lower priority than the other

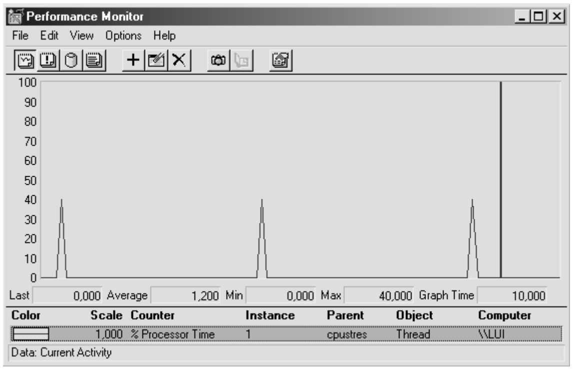

Last 0,000 Average 1,200 Min 0,000 Max 40,000 Graph Time 10,000

Color	Scale	Counter	Instance	Parent	Object	Computer
	1,000	% Processor Time	1	cpustres	Thread	\\LUI

Data: Current Activity