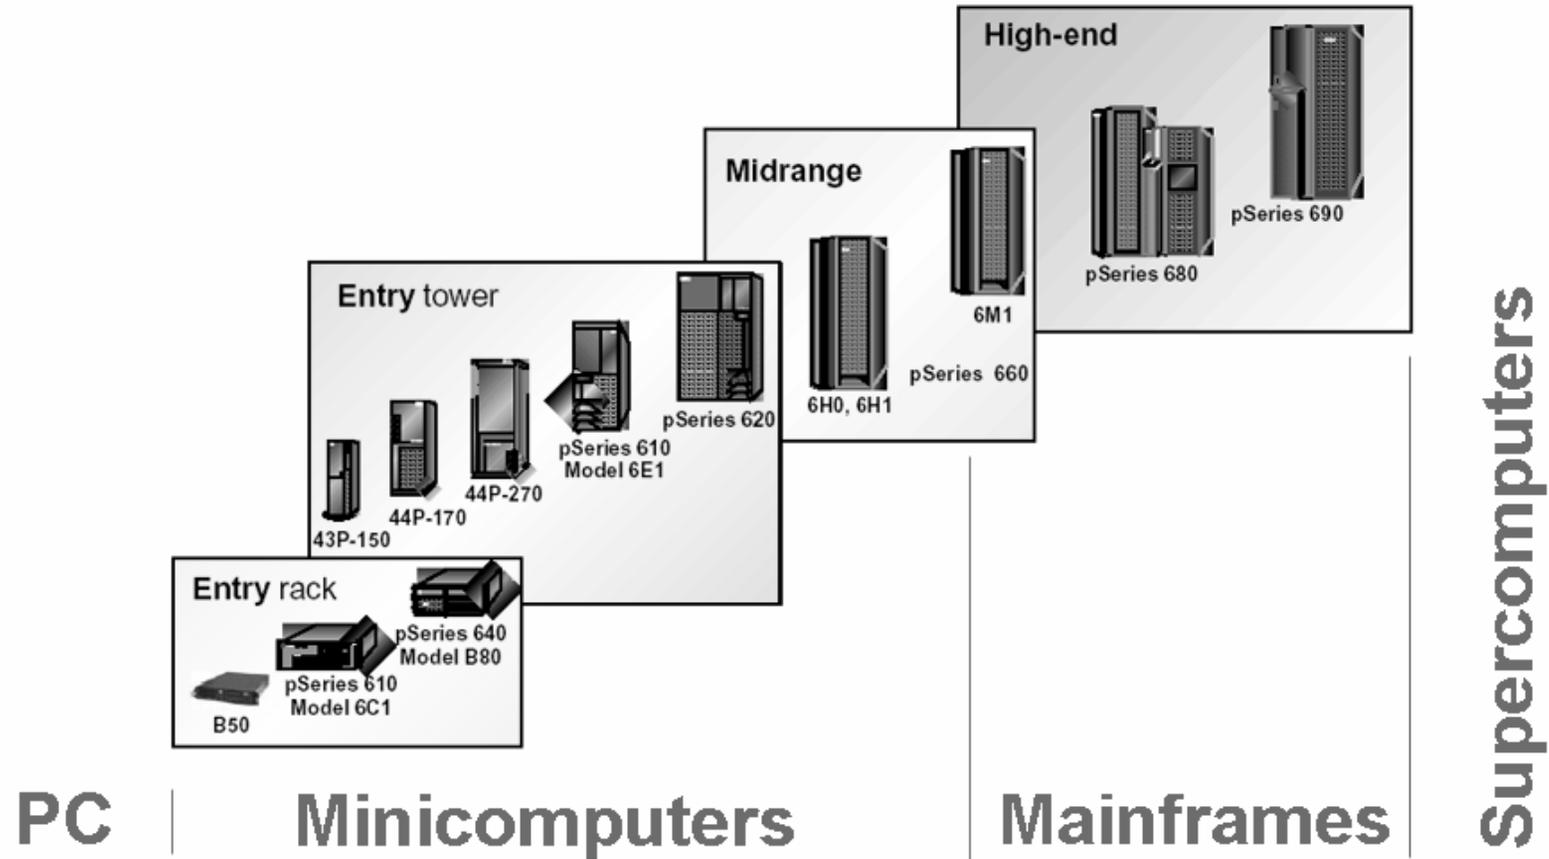


Introduction to Operating Systems

Dr. Zulfikar H.A. Kassam

Types of Computers



Objectives for this course

COMPUTER SYSTEM

Identify and define components of a computer system.

OPERATING SYSTEMS (Part 1)

Identify and define components of an Operating System (OS).

Explain memory management.

Explain processor management.

OPERATING SYSTEMS (Part 2)

Explain device management.

Explain file management.

INTRODUCTION TO COMPUTERS (HARDWARE)

Basic Definitions

Bits & Bytes

- ◆ Current Computers are digital
- ◆ All the data is binary
- ◆ 2 possible values only, 0 or 1
- ◆ Each piece of binary data is defined as a bit
- ◆ A sequence of 8 bits called a byte.
- ◆ The bit is referred to by a lowercase "b"
- ◆ Network cards speeds are designated in bits -
e.g. 100Mbps (100 Mega bits per second = $100/8$
(12.5) Mega Bytes per second)
- ◆ The Byte is referred to by an uppercase "B"
- ◆ Example: Memory (RAM) is 64MB (64 Mega Bytes)

Word

- ◆ The number of bits or bytes that a computer can process together is called a WORD.
- ◆ E.g. Intel Pentium processors are 32-bit systems (1 Word = 4 Bytes = 32 bits)
- ◆ 32 bits are processed in 1 clock cycle.
- ◆ E.g. Sun Microsystems and IBM have 64 bit architecture (1 Word - 8 Bytes = 64 bits)
- ◆ 64 bits are processed in 1 clock cycle
- ◆ Roughly, 600MHz 64-bit machine = 1.2 GHz 32-bit machine
- ◆ Intel Itanium processors are 64-bit machines running at 700-900 MHz.

Central Processing Unit (CPU)

- ◆ CPU controls all the functions performed by a computer
- ◆ The CPU and its architecture (32 bit or 64 bit) are the primary components that determine the speed of the computer.

Memory

- ◆ Although memory is technically any form of electronic storage, it is used most often to identify fast, temporary forms of storage.
- ◆ If your computer's CPU had to constantly access the hard drive to retrieve every piece of data it needs, it would operate very slowly.
- ◆ Typically, Memory is 1000 times faster than Hard Drive
- ◆ When the information is kept in memory, the CPU can access it much more quickly.

Types of Memory

The information can be stored either in the

- ◆ ROM (Read Only Memory)

- Non-volatile memory that does not lose its data when the system or device is turned off

- ◆ RAM (Random Access Memory)

- Volatile memory that loses any data as soon as the system is turned off; it requires constant power to retain memory contents.

ROM

- ◆ ROM - Read Only Memory
- ◆ PROM - Programmable Read Only Memory (uses Ultra-Violet Rays)
- ◆ EPROM - Erasable Programmable Read Only Memory (uses Ultra-Violet Rays)
- ◆ EEPROM (also known as Flash Memory) - Electrically Erasable Programmable Read Only Memory
- ◆ BIOS - uses EEPROM

BIOS

(Basic Input/Output System)

- ◆ On computer startup, the computer loads the **BIOS** from the ROM
- ◆ The BIOS provides
 - Date and time
 - the most basic information about storage devices (Floppy drive, Hard drive, CDROM, ZIP/JAZZ drives)
 - boot sequence (which storage device boots up first)
 - security
 - **Plug and Play** (auto device recognition)
 - and few other items

Types of RAM

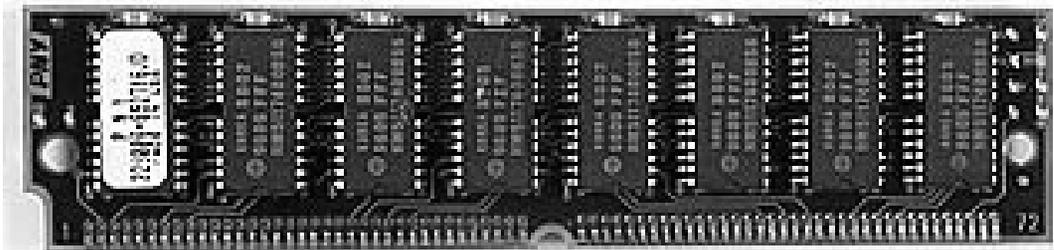
- ◆ SRAM - Static RAM
 - Only used for Cache Memory because it is faster (5-10 times faster than RAM) and more expensive
- ◆ DRAM - Dynamic RAM
 - Typical RAM used in computers

Types of DRAM

- ◆ EDO DRAM - Extended Data-Out Dynamic Random Access Memory
- ◆ SDRAM - Synchronous Dynamic Random Access Memory (133 MHz)
- ◆ DDR DRAM - Double Data Rate (266 MHz) DRAM
- ◆ RDRAM - Rambus Dynamic Random Access Memory

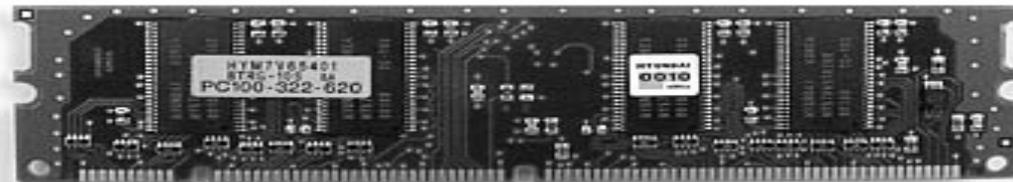
SIMMs

- ◆ DRAM and EDO RAM used to come in 72 pin format
- ◆ Used to be called SIMMs
- ◆ Single In-line Memory Module
- ◆ Had to be used in pairs



DIMMs

- ◆ EDO RAM, SDRAM, DDR DRAM, RDRAM come in 168 pin or 184 pin format
- ◆ They are called DIMMs
- ◆ Dual In-line Memory Module
- ◆ Can be used in singles except for some DDR DRAM has to be used in pairs, and RDRAM which have to be used in groups of 4



Storage Devices for PCs

- ◆ Ejectable - Floppy, CD ROM, DVD ROM, ZIP, JAZZ
- ◆ Fixed - IDE/ATA Hard drives, SCSI Hard Drives

IDE/ATA Drives

- ◆ IDE - Integrated Drive Electronics (IDE) interface
- ◆ ATA - AT Attachment (ATA) to reflect attachment to IBM AT computers
- ◆ Popular because they are cheap
- ◆ 4 devices maximum
 - Primary Master, Primary Slave, Secondary Master, Secondary Slave
- ◆ Maximum size ~ 200GB



SCSI (Small Computer System Interface) Drives

Advantages

- ◆ It's fast -- up to 320 Mega Bytes per second (MBps).
- ◆ It's very reliable.
- ◆ It allows you to put more devices on one bus (total 14).
- ◆ It works on most computer systems (UNIX, Linux, Windows, and Apple).
- ◆ Maximum size ~ 144GB

How do we make the capacity of the storage bigger?

- ◆ Create a RAID array
- ◆ What is RAID?
- ◆ RAID stands for Redundant Array of Inexpensive Disks
- ◆ Large storage capacities can be created by using RAID 5
- ◆ RAID 5 - means Disk Striping with Parity

RAID variations

- ◆ RAID is available for IDE/ATA drives, SCSI drives and Fibre Channel drives
- ◆ Small businesses tend to use IDE/ATA due to being inexpensive
- ◆ Medium to Large enterprises tend to use SCSI or Fibre Channel due to better performance and reliability
- ◆ Fibre Channel is the best with transfer rates of 2 Gb/s each way (up to 500MB/s) as compared to SCSI (320MB/s) and IDE (30MB/s)

Mass Storage Devices (use RAID)

- ◆ Network Attached Storage (NAS)
- ◆ Manufactured by Network Appliance and EMC
- ◆ Storage Area Networks (SAN)
- ◆ Manufactured by EMC, IBM, Hitachi
- ◆ High Quality Mass Storage Devices use only SCSI or Fibre Channel (faster than SCSI)

Tape Backup

- ◆ Why use Tape Backup?
- ◆ Problem: Hard drive data may occasionally be corrupted, Unintentional change
- ◆ Solution: We need to backup the information from the hard drive into tapes using a tape backup system.
- ◆ Most Operating Systems have simple in-built Backup software
- ◆ Disadvantage: In-built Backup software cannot backup open files.

Tape Backup (contd.)

Advantages

- ◆ Relatively inexpensive
- ◆ Allow large amount of information (up to 100 GB) to be stored in a small compact tape
- ◆ It is removable media allowing easy transportation of large amounts of information - useful in the case of disaster such as flooding or earthquakes

Disadvantages

- ◆ Media is fragile (the tape can break from repeated usage)
- ◆ Data transfer rates are extremely slow (few MBs per minute).

Input/Output Devices

- ◆ The computer need many other peripherals, typically called I/O devices to be functional
- ◆ I/O devices are needed to input or output information to and from the CPU.
- ◆ There are many different types of I/O devices that a convention computer needs.

Input/Output Devices

- ◆ Video/Graphics Card
- ◆ Keyboard
- ◆ Mouse
- ◆ Serial Ports
- ◆ Parallel Port
- ◆ USB Ports

Free BSD UNIX

<ftp://ftp.freebsd.org/pub/FreeBSD/releases/i386/ISO-IMAGES/4.4>

http://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/index.html

RedHat Linux 7.3

- ◆ <ftp://ftp.redhat.com/pub/redhat/linux/7.3/en/iso/i386/>

How a Computer Works

- ◆ You turn the computer on.
- ◆ The computer loads data from read-only memory (ROM)
- ◆ Computer performs a Power-On Self-Test (POST) to make sure all the major components are functioning properly.
- ◆ The computer loads the basic input/output system (BIOS) from ROM.

How a Computer Works (2)

- ◆ The BIOS provides the most basic information about storage devices, boot sequence, security, Plug and Play (auto device recognition) capability and a few other items.
- ◆ The computer loads the operating system (OS) from the hard drive into the system's RAM.

How a Computer Works (3)

- ◆ Important! Only the critical parts of the operating system are maintained in RAM. This allows the CPU to have fast access to the operating system, which enhances the performance and functionality of the overall system.
- ◆ When an application is opened, it is loaded into RAM.
- ◆ To conserve RAM usage, many applications load only the essential parts of the program initially and then load other pieces as needed.

How a Computer Works (4)

- ◆ After an application is loaded, any files that are opened for use in that application are loaded into RAM.
- ◆ When you save a file and close the application, the file is written to the specified storage device, and then the file and the application are purged from RAM.
- ◆ Problem: Memory Leaks occur if the memory occupied by the application is not cleared completely.

OPERATING SYSTEMS

Discussion Questions

- ◆ What is an Operating System?
- ◆ What kinds of Operating Systems are used currently?
- ◆ Based on your experiences, what are the Advantages and Disadvantages of these Operating Systems?
- ◆ Why is Windows very popular?
- ◆ What would you like to see improve?

Operating System (Definition)

- ◆ An Operating System is a system software that interfaces between the application programs and the hardware and manages the resources of the overall computer system
- ◆ Facilitates the building of Application Software with minimal knowledge of the hardware

Examples of Operating Systems

- ◆ Windows 95, Windows 98, Windows NT, Windows 2000, Windows ME, Windows XP, Windows 2003
- ◆ UNIX, Linux, Irix, and all other members of the UNIX family
- ◆ MacOS 7, MacOS 8, Mac OS 9, Mac OS X and all other members of the MacOS family

Types of Operating Systems

- ◆ Real-time operating system (RTOS)
- ◆ Single-user, single task
- ◆ Single-user, multi-tasking
- ◆ Multi-user, multi-tasking

Real-Time Operating System

- ◆ Real-time operating systems are used to control machinery, scientific instruments and industrial systems.
- ◆ An RTOS typically will be a "sealed box" when delivered for use.
- ◆ A very important part of an RTOS is managing the resources of the computer so that a particular operation executes in precisely the same amount of time every time it occurs.
- ◆ In a complex machine, having a part move more quickly just because system resources are available may be just as catastrophic as having it not move at all because the system is busy.

Single-user, single task

- ◆ this operating system is designed to manage the computer so that one user can effectively do one thing at a time.
- ◆ E.g. The Palm OS for Palm handheld computers is a good example of a modern single-user, single-task operating system.
- ◆ E.g. MS-DOS, PC-DOS

Single-user, multi-tasking

- ◆ This is the type of operating system most people use on their desktop and laptop computers today.
- ◆ Windows 98, MacOS, Windows NT are all examples of an operating system that will let a single user have several programs in operation at the same time.
- ◆ For example, it's entirely possible for a Windows user to be writing a note in a word processor while downloading a file from the Internet while printing the text of an e-mail message and sharing their files and printers.

Multi-User, Multi-Tasking

- ◆ Many different users can log onto one computer to take advantage of the computer's resources simultaneously.
- ◆ The OS must make sure that the requirements of the various users are balanced
- ◆ Each of the programs they are using has sufficient and separate resources so that a problem with one user doesn't affect the entire community of users.
- ◆ Unix, Linux, Windows XP, Windows NT Terminal Server, and Windows 2000 Terminal Server (with or without Citrix) and mainframe operating systems, such as MVS, are examples of multi-user operating systems.

Implications of an Operating Systems

Problem:

- ◆ Operating System will be Hardware dependent

Solution: (implemented by Microsoft in Windows NT Architecture)

- ◆ Not allow the major components of the Operating System to talk directly to hardware. Has to talk through the Hardware Abstraction Layer (HAL).

Hardware Abstraction Layer in NT

Advantages

- ◆ Only a small part of the code had to be rewritten, i.e., the HAL.
- ◆ Windows NT 4.0 was supported on Intel, DEC Alpha, MIBS, and PowerPC.

Disadvantage

- ◆ Many PC Games do not work on NT because the games try to directly access the Video Card and Sound Card.

Processor and Memory Management

Discussion

- ◆ A company has 50 Windows 98 desktop computers running Office 97, 1 Exchange Server for Email, 1 IIS Server (Web Server) and 2 NT Domain Controllers. They are thinking of upgrading. You have been hired as a consultant to determine whether it is more beneficial to move to Windows 2003 based Servers and Windows XP desktops with Office XP 2003 or move to Free Linux Servers and Workstations and buy only StarOffice suite from Sun.
- ◆ Determine all the information that you need to accumulate to allow you to determine which solution you will choose.

Current Hardware

- ◆ Servers: Pentium III-733, 256MB RAM, 18GB SCSI
- ◆ Workstations: Pentium II-300, 64MB RAM, 10GB HD

Components of an Operating System

- ◆ Processor management.
- ◆ Memory management.
- ◆ Storage management (file system management).
- ◆ I/O device management.

Processor Management

- ◆ Ensure that each process and application receives enough of the processor's time to function properly.
- ◆ It uses Scheduling and Queuing techniques to achieve this goal.
- ◆ The basic unit of software that the operating system deals with for scheduling the work done by the processor is a **process thread**.

Processes

- ◆ A **process** is an executable program.
- ◆ Executable program may be an Application program or a service/daemon.
- ◆ E.g. Winword is a Application process that runs when Microsoft Word is launched.
- ◆ E.g. Spoolsv.exe, snmp.exe are services/daemons

Image Name	PID	CPU	CPU Time	Mem Usage
WINWORD.EXE	3144	00	0:16:55	17,120 K
WinMgmt.exe	1428	00	0:00:20	144 K
WINLOGON.EXE	228	00	0:00:23	1,280 K
VsTskMgr.exe	840	00	0:00:29	2,776 K
tftpd.exe	1396	00	0:00:02	2,612 K
termsrv.exe	1356	00	0:00:20	3,036 K
tcpsvcs.exe	1232	00	0:00:01	2,704 K
taskmgr.exe	3552	03	0:00:02	1,144 K
System Idle Process	0	96	156:37:28	16 K
System	8	00	0:08:48	212 K
svchost.exe	1320	00	0:00:01	2,632 K
svchost.exe	648	00	0:00:11	6,676 K
svchost.exe	448	00	0:00:07	3,440 K
stisvc.exe	1284	00	0:00:03	1,612 K
SPOOLSV.EXE	488	00	0:00:10	2,572 K
snmp.exe	1276	00	0:00:14	2,868 K
SMSS.EXE	176	00	0:00:01	348 K
shstat.exe	2216	00	0:00:14	1,992 K

Show processes from all users

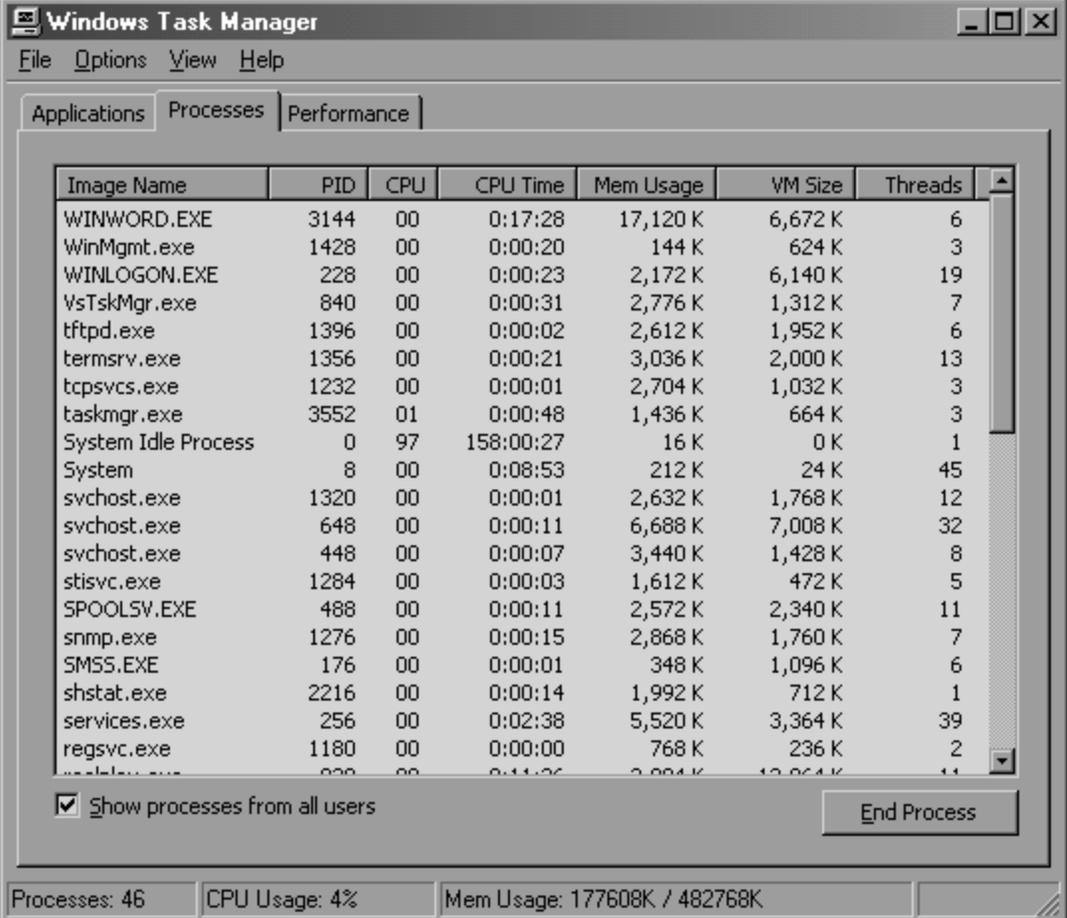
End Process

Processes: 46 CPU Usage: 4% Mem Usage: 175504K / 482768K

- All process have a process ID

Threads

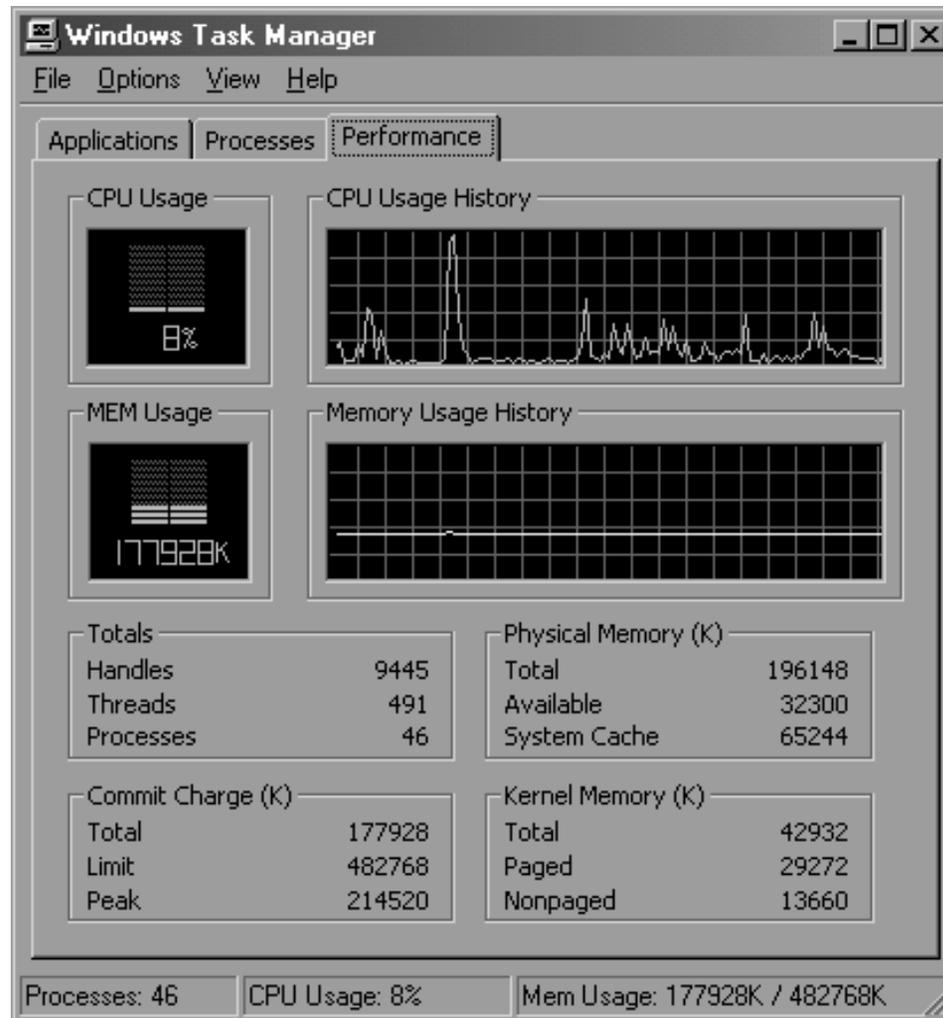
- ◆ A **thread** is a set of commands within a process to which the OS grants processor time.
- ◆ A process must have at least one thread.
- ◆ A process having more than one thread is called a multi-threaded process (or program/application)



The screenshot shows the Windows Task Manager window with the 'Processes' tab selected. The window title is 'Windows Task Manager' and it has a menu bar with 'File', 'Options', 'View', and 'Help'. Below the menu bar are three tabs: 'Applications', 'Processes', and 'Performance'. The 'Processes' tab is active, displaying a table of running processes. The table has columns for 'Image Name', 'PID', 'CPU', 'CPU Time', 'Mem Usage', 'VM Size', and 'Threads'. The 'Threads' column shows the number of threads for each process. At the bottom of the window, there is a status bar with 'Processes: 46', 'CPU Usage: 4%', and 'Mem Usage: 177608K / 482768K'. There is also a checkbox for 'Show processes from all users' and an 'End Process' button.

Image Name	PID	CPU	CPU Time	Mem Usage	VM Size	Threads
WINWORD.EXE	3144	00	0:17:28	17,120 K	6,672 K	6
WinMgmt.exe	1428	00	0:00:20	144 K	624 K	3
WINLOGON.EXE	228	00	0:00:23	2,172 K	6,140 K	19
VsTskMgr.exe	840	00	0:00:31	2,776 K	1,312 K	7
tftpd.exe	1396	00	0:00:02	2,612 K	1,952 K	6
termsrv.exe	1356	00	0:00:21	3,036 K	2,000 K	13
tcpsvcs.exe	1232	00	0:00:01	2,704 K	1,032 K	3
taskmgr.exe	3552	01	0:00:48	1,436 K	664 K	3
System Idle Process	0	97	158:00:27	16 K	0 K	1
System	8	00	0:08:53	212 K	24 K	45
svchost.exe	1320	00	0:00:01	2,632 K	1,768 K	12
svchost.exe	648	00	0:00:11	6,688 K	7,008 K	32
svchost.exe	448	00	0:00:07	3,440 K	1,428 K	8
stisvc.exe	1284	00	0:00:03	1,612 K	472 K	5
SPOOLSV.EXE	488	00	0:00:11	2,572 K	2,340 K	11
snmp.exe	1276	00	0:00:15	2,868 K	1,760 K	7
SMSS.EXE	176	00	0:00:01	348 K	1,096 K	6
shstat.exe	2216	00	0:00:14	1,992 K	712 K	1
services.exe	256	00	0:02:38	5,520 K	3,364 K	39
regsvc.exe	1180	00	0:00:00	768 K	236 K	2
...

Processes & Threads



Multithreading

- ◆ A process having more than one thread is called a multi-threaded process (or program/application)

E.g Internet Explorer Web Browser

- ◆ As you type the URL of a Website,
- ◆ One process thread is accepting input from the keyboard
- ◆ Another process thread is searching your browsing history to see if the typed info matches any sites you recently visited.
- ◆ If this second thread discovers a match, it will display this match

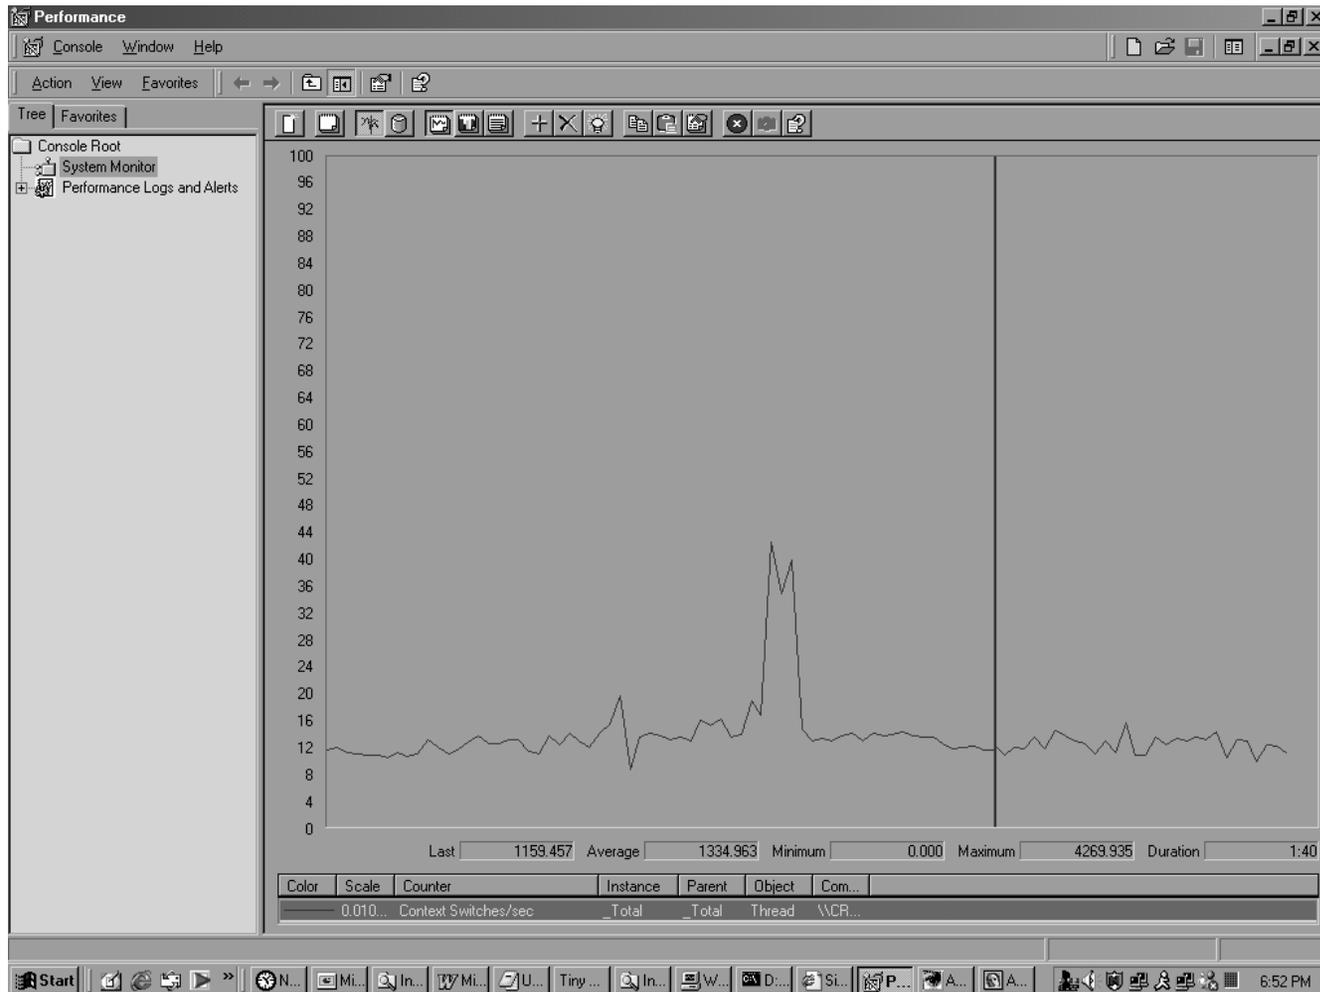
Multitasking

- ◆ Ability of an Operating System to run multiple applications at the same time
- ◆ Previous versions of Windows used **Cooperative Multitasking** - each application had to be “polite” and give a chance for other applications to run - caused a lot of system hangups
- ◆ Windows NT, 2000 and XP use **Pre-emptive Multitasking** - the processor will allow a thread to run for a predefined time and suspend this thread, and then execute another thread - minimizes system hangups

Multitasking - How it's done

- ◆ A single CPU can only do one thing at a time
- ◆ The operating system must arrange the execution of applications so that you believe that there are several things happening at once
- ◆ the operating system has to switch between different process threads thousands of times a second - this process is known as **context-switching**
- ◆ the operating system allots a certain number of CPU execution cycles for a thread.

Thread-Context Switching



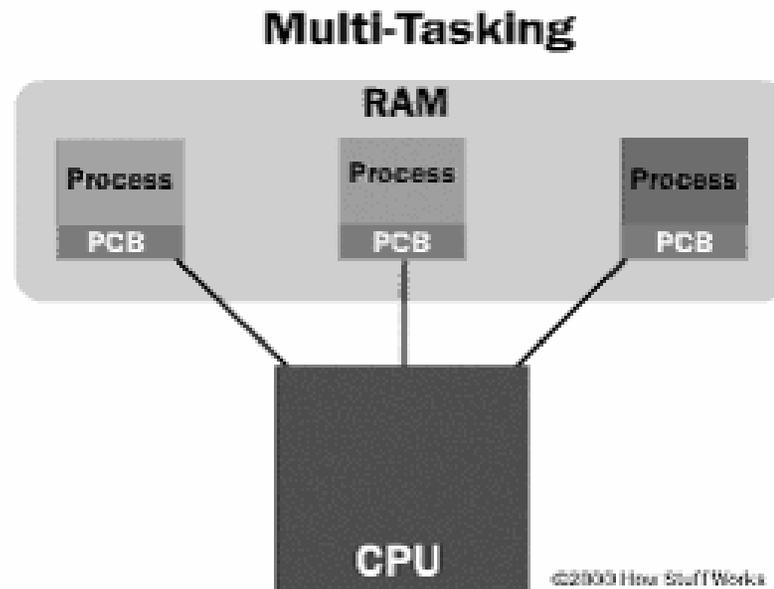
Started Adobe Photoshop and Illustrator -
Context Switching = 1200 to 4000

Multitasking - How it's done *(continued)*

- ◆ After that number of cycles, the operating system makes copies of all the registers, stacks and queues used by the process thread, and notes the point at which the process thread paused in its execution.
- ◆ It then loads all the registers, stacks and queues used by the second process thread and allows it a certain number of CPU cycles.
- ◆ When those are complete, it makes copies of all the registers, stacks and queues used by the second process thread, and loads the first process thread.

Process Control Block

- ◆ All of the information needed to keep track of a process when switching is kept in a data package called a **process control block**.



Symmetric and Asymmetric Multiprocessing

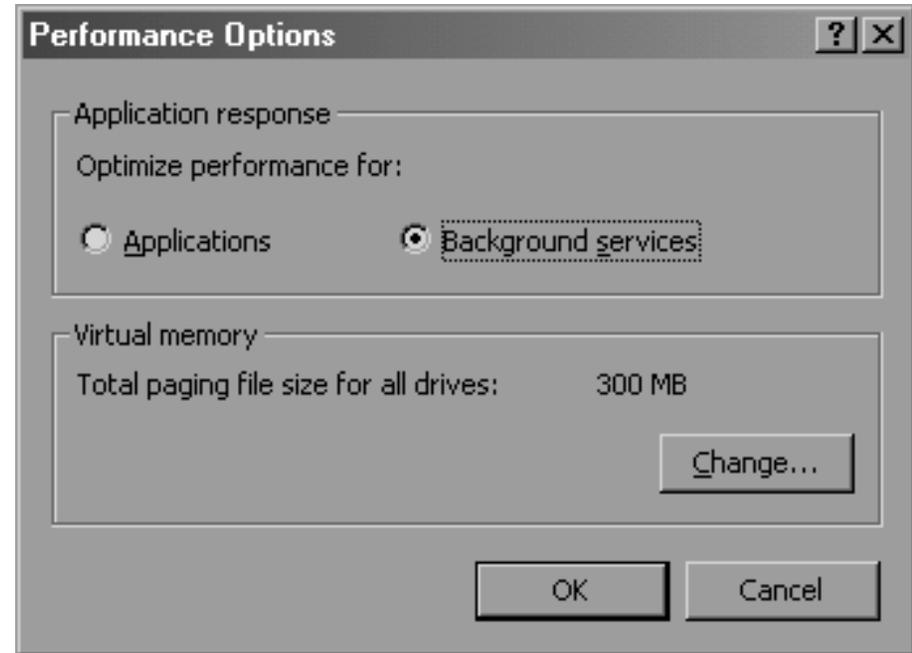
- ◆ The ability of a system to utilize more than one processor is called Multiprocessing.
- ◆ In a system with two or more CPUs, the operating system must divide the workload among the CPUs
- ◆ **Asymmetric Multiprocessing** operating systems dedicate one or more CPUs for the Operating System alone and divide application processes among the remaining CPUs.
- ◆ **Symmetric Multiprocessing** operating systems divide the Operating System and Application processes equally between the CPUs
- ◆ Windows NT and Windows 2000 use Symmetric Multiprocessing (SMP)

Dispatcher

- ◆ Consider 2 programs running concurrently that are in a **ready state**
- ◆ **Ready state** implies ready to be executed
- ◆ Which one should go first?
- ◆ The decision is made by an operating system routine called the **dispatcher**.
- ◆ The dispatcher will execute the program with the highest priority.
- ◆ How to determine priority?

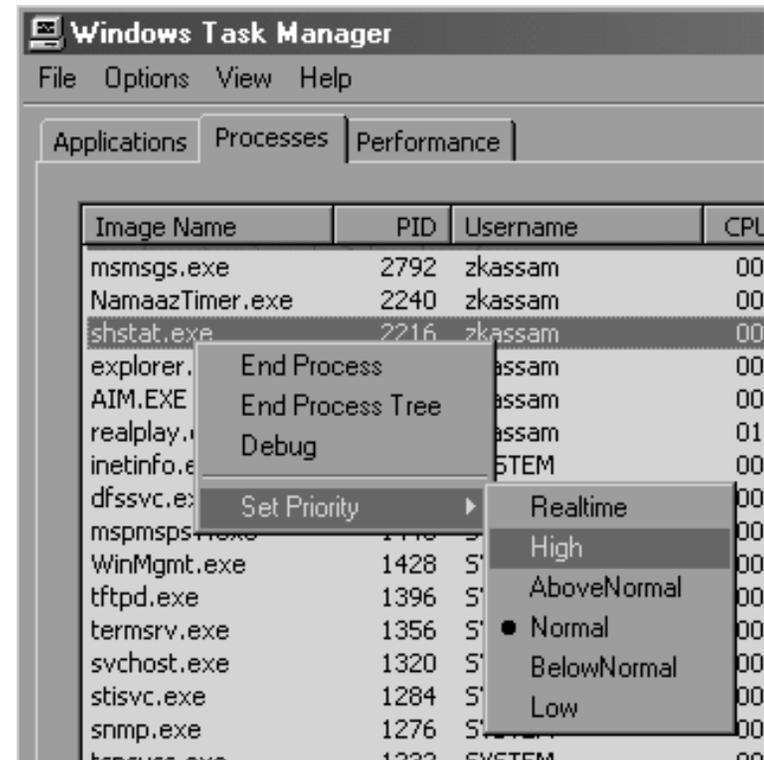
Foreground and Background Processes

- ◆ Applications (like Word, Excel) are usually foreground processes
- ◆ Services (Spooler, SNMP, Server, WWW, FTP) are background processes
- ◆ Which has priority will depend on the configuration settings



Process Priority

- ◆ Priority of a process can be further determined assigning it higher or lower priority.
- ◆ Only processes running under the Username context can have their priority changed



- ◆ `START [/LOW | /NORMAL | /HIGH | /REALTIME | /ABOVENORMAL | /BELOWNORMAL] [command/program]`

Interrupts

- ◆ A process thread being executed by the processor may be temporarily suspended from execution if the processor receives an interrupt from an Input/Output device.
- ◆ E.g., Movements of the mouse and keyboard input are the most common form of interrupts that a CPU receives
- ◆ Sometimes the operating system will schedule the priority of processes so that interrupts are **masked** -- that is, the operating system will wait for a running thread to complete execution before executing the interrupts from some sources so that a particular job can be finished as quickly as possible. These are called **maskable interrupts**.
- ◆ There are some interrupts (such as those from error conditions or problems with memory) that are so important that they need to be dealt with immediately, regardless of the other tasks at hand. These are called **non-maskable interrupts** (NMIs).

User and Kernel Modes

- ◆ The components that run in **User Mode** do not directly access the hardware resources.
- ◆ All applications run in User Mode
- ◆ Implications: Application running in User Mode can crash without de-stabilizing the overall operation.
- ◆ E.g., Word can crash without other applications being affected.

User and Kernel Modes *(continued)*

- ◆ The components that run in **Kernel Mode** directly access the hardware resources (via the Hardware Abstraction Layer in NT).
- ◆ The File System Manager, Plug and Play Manager, Device Drivers, Virtual Memory Manager run in Kernel Mode
- ◆ Implications: If any component running in Kernel Mode crashes, it can de-stabilizing the overall operation.
- ◆ E.g., a poor device driver for your network card or SCSI card can crash your system

Command Interpreter or Shell

- ◆ interface between user and OS
- ◆ used to transform a request from the user into a request to the OS
- ◆ can be GUI (graphical user interface) or line-oriented
- ◆ the appearance of the command interpreter is the principal feature of the OS noted by users
- ◆ In Windows, the command interpreter is based on a graphical user interface
- ◆ In UNIX, there is a line-oriented command interpreter - KDE and GNOME offer GUI
- ◆ In Linux, KDE and GNOME also offer GUI

UNIX Systemboard

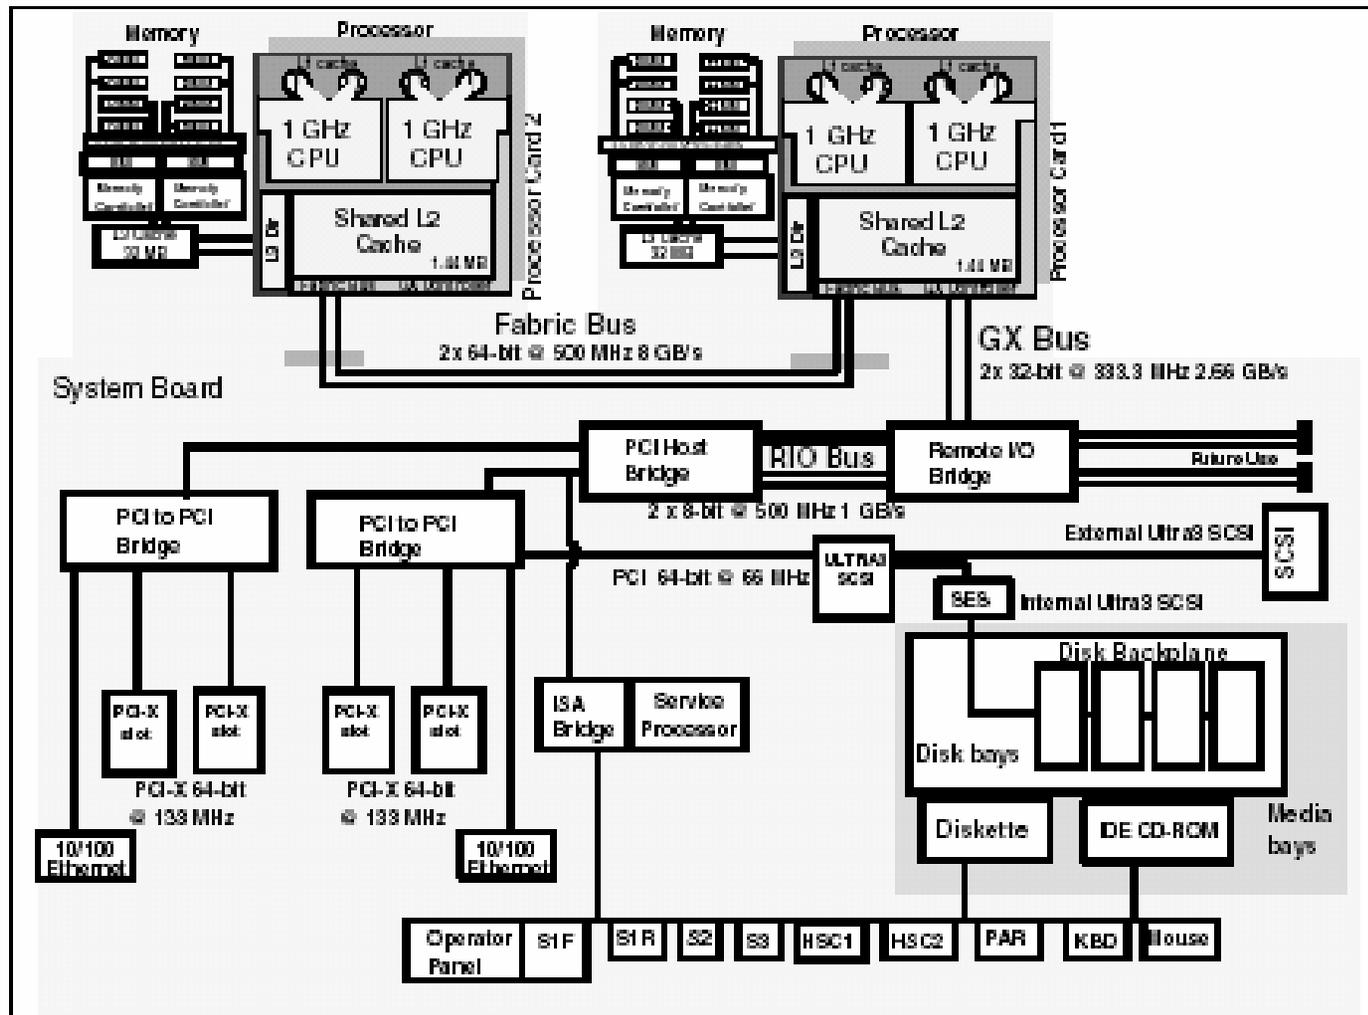
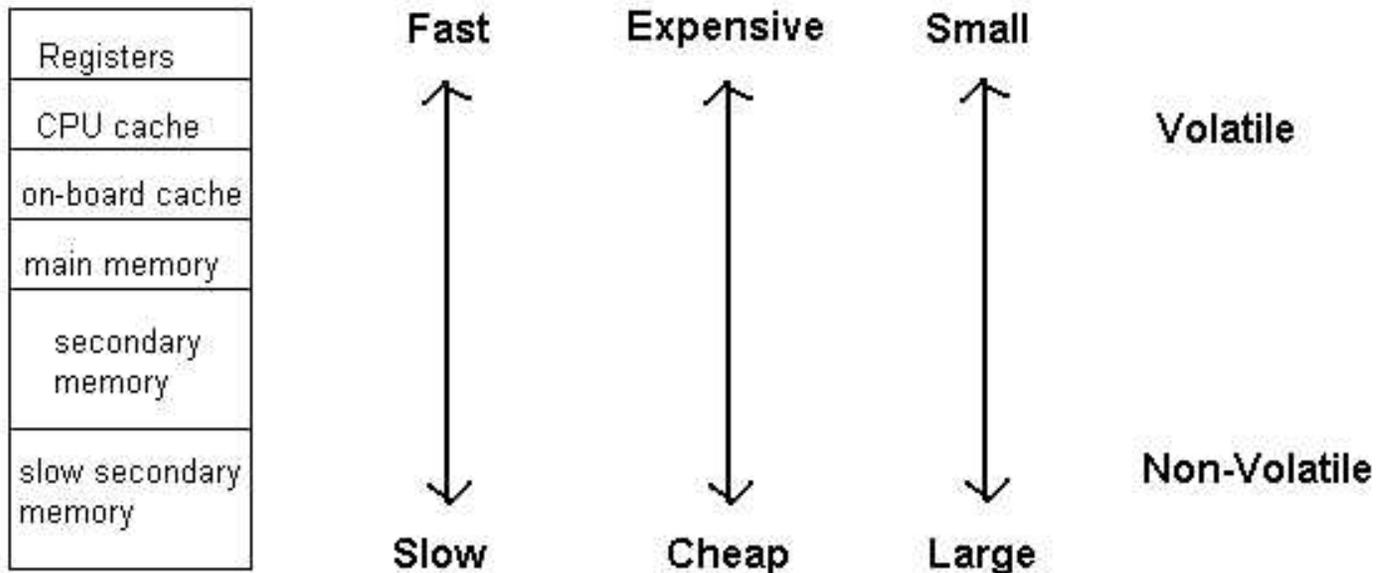


Figure 2-1 Conceptual diagram of the Models 6C4 and 6E4 system architecture

Memory Management



Memory Management *(continued)*

- ◆ CPU Cache (L1 Cache) - memory located on the processor chip (VOLATILE)
- ◆ On-board Cache (L2 Cache) - located on circuit board or part of a processor module/board; fastest external memory available (VOLATILE)
- ◆ Main Memory (RAM) - i.e., SIMMS, DIMMS, (VOLATILE)
- ◆ Secondary Memory - i.e., hard drive (NON-VOLATILE)
- ◆ Slow secondary memory - e.g., tapes, diskettes (NON-VOLATILE)

Volatile and Non-Volatile Memory

Volatile Storage

- ◆ contents lost if power is interrupted

Non-volatile Storage

- ◆ can withstand power failures and system crashes

Cache

- ◆ A cache is a very fast block of memory that speeds up the performance of another device.
- ◆ Frequently used data are stored in the cache.
- ◆ The computer looks in the cache first to see if what it needs is there.
- ◆ Level 1 Cache ($\sim 32\text{K}$) is located directly inside the CPU itself and has the most direct effect on overall performance (bus speed = processor speed).
- ◆ Level 2 Cache ($\sim 256\text{K}-512\text{K}$) is located on the motherboard or processor module/board. It stores frequently used data from the computer's main memory (RAM). (bus speed may be $1/4$ to $1/2$ of processor speed)
- ◆ Pentium Pro had Level 2 Cache on the processor - as fast as Level 1 Cache)

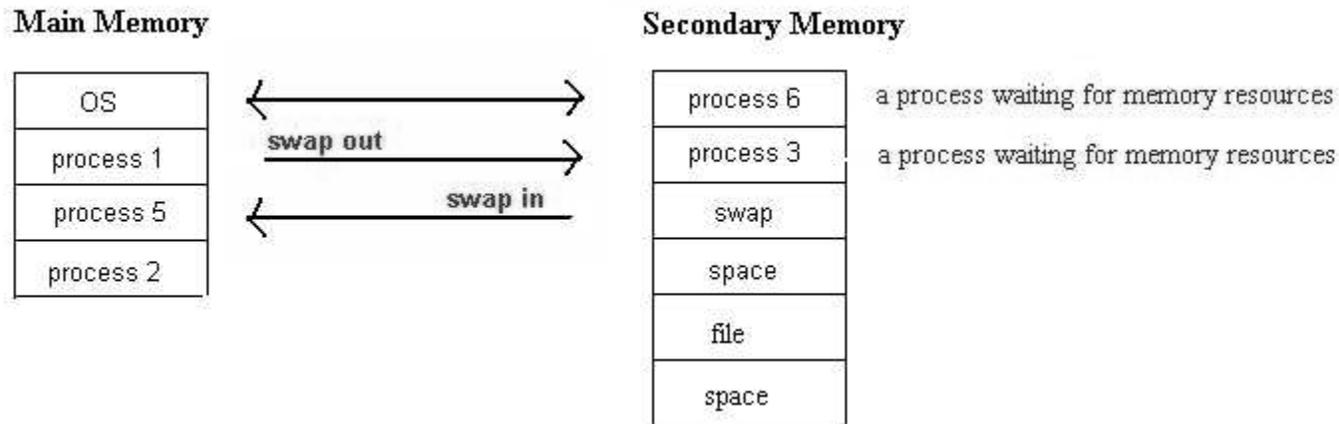
Memory Management - Introduction

- ◆ Processes must be brought into RAM before it can be executed.
- ◆ RAM does not usually have adequate capacity to hold all the processes
- ◆ Therefore, it stores only the minimum required processes in RAM
- ◆ The rest is stored in Virtual Memory (VM)
- ◆ The VM contents are stored in a Page File on the hard drive

Virtual Memory

- ◆ Virtual memory – separation of user logical memory from physical memory.
- ◆ Only part of the process needs to be in memory for execution.
- ◆ Logical address space can therefore be much larger than physical address space.
- ◆ Need to allow memory contents to be swapped in and out.

Swapping



- ◆ Swapping - moves partial or entire processes between main (RAM) and secondary memory (Virtual Memory on Hard Drive)

Swapping

- ◆ A process can be swapped temporarily out of memory (RAM) to a backing store (Virtual Memory), and then brought back into memory for continued execution.
- ◆ Backing store – fast disk large enough to accommodate copies of all memory images for all users; must provide direct access to these memory images.
- ◆ Major part of swap time is transfer time; total transfer time is directly proportional to the amount of memory swapped.
- ◆ Modified versions of swapping are found on many systems, i.e., UNIX and Microsoft Windows.

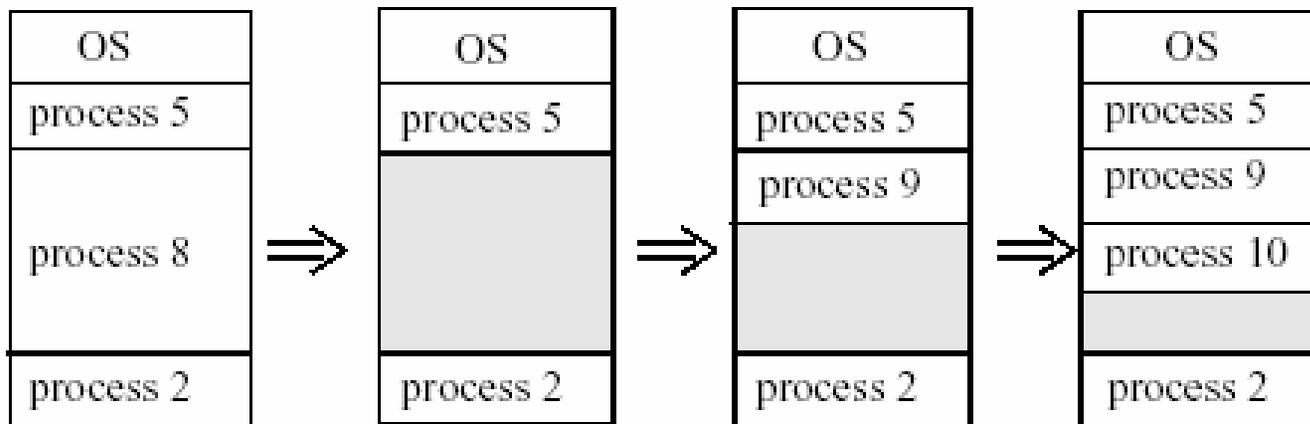
Roll Out / Roll In

- ◆ Roll out, roll in – variation of swapping technique used for priority-based scheduling algorithms
- ◆ lower-priority process is swapped out so that higher-priority process can be loaded and executed.

Contiguous Allocation

- ◆ Hole– block of available memory; holes of various size are scattered throughout memory.
- ◆ When a process arrives, it is allocated memory from a hole large enough to accommodate it.

Operating system maintains information about:
a) allocated regions b) free regions (hole)



Demand Paging

- ◆ Bring a page into memory only when it is needed.
- ◆ Less I/O needed
- ◆ Less memory needed
- ◆ Faster response
- ◆ More users
- ◆ Page is needed -> reference to it
 - not-in-memory -> bring to memory

Swapping (paging) Example

page 0
page 1
page 2
page 3
page 4
page 5
page 6

logical memory space
(virtual address space)

logical memory space	physical memory space
page 0	page 1
page 2	page 3
page 4	page 2
page 5	page 0

virtual memory manager
page table

Frame	Contents
0	page 5
1	page 0
2	page 4
3	page 2
4	-

physical
memory

page 1
page 3
page 6

page file
(stored in the
hard drive)

Paging

- ◆ Divide physical memory into fixed-sized blocks called frames (size between 512 bytes and 8192 bytes).
- ◆ Divide logical memory into blocks of same size called pages.
- ◆ E.g Windows NT and Windows 2000 use 4K page sizes
- ◆ Keep track of all free frames.
- ◆ To run a program of size n pages, need to find n free frames and load program.
- ◆ Set up a page table to translate logical to physical addresses.
- ◆ *Problem:* Internal fragmentation of the Memory as well as page file - leads to poor performance.

Page Fault

- ◆ When virtual address points to a page that is not in physical memory, a page fault is recognized.
- ◆ A page-in (or swap-in) operation begins.
- ◆ Monitoring the page faults/sec and pages/sec provides good indication if memory is insufficient

Thrashing

- ◆ When physical memory (RAM) is full, a demand for new page means another page must be swapped out first.
- ◆ If this happens frequently, system may end up spending so much time swapping pages into and out of memory that the computer response becomes extremely slow while the hard drive runs vigorously.

Network Devices Management

What is network management?

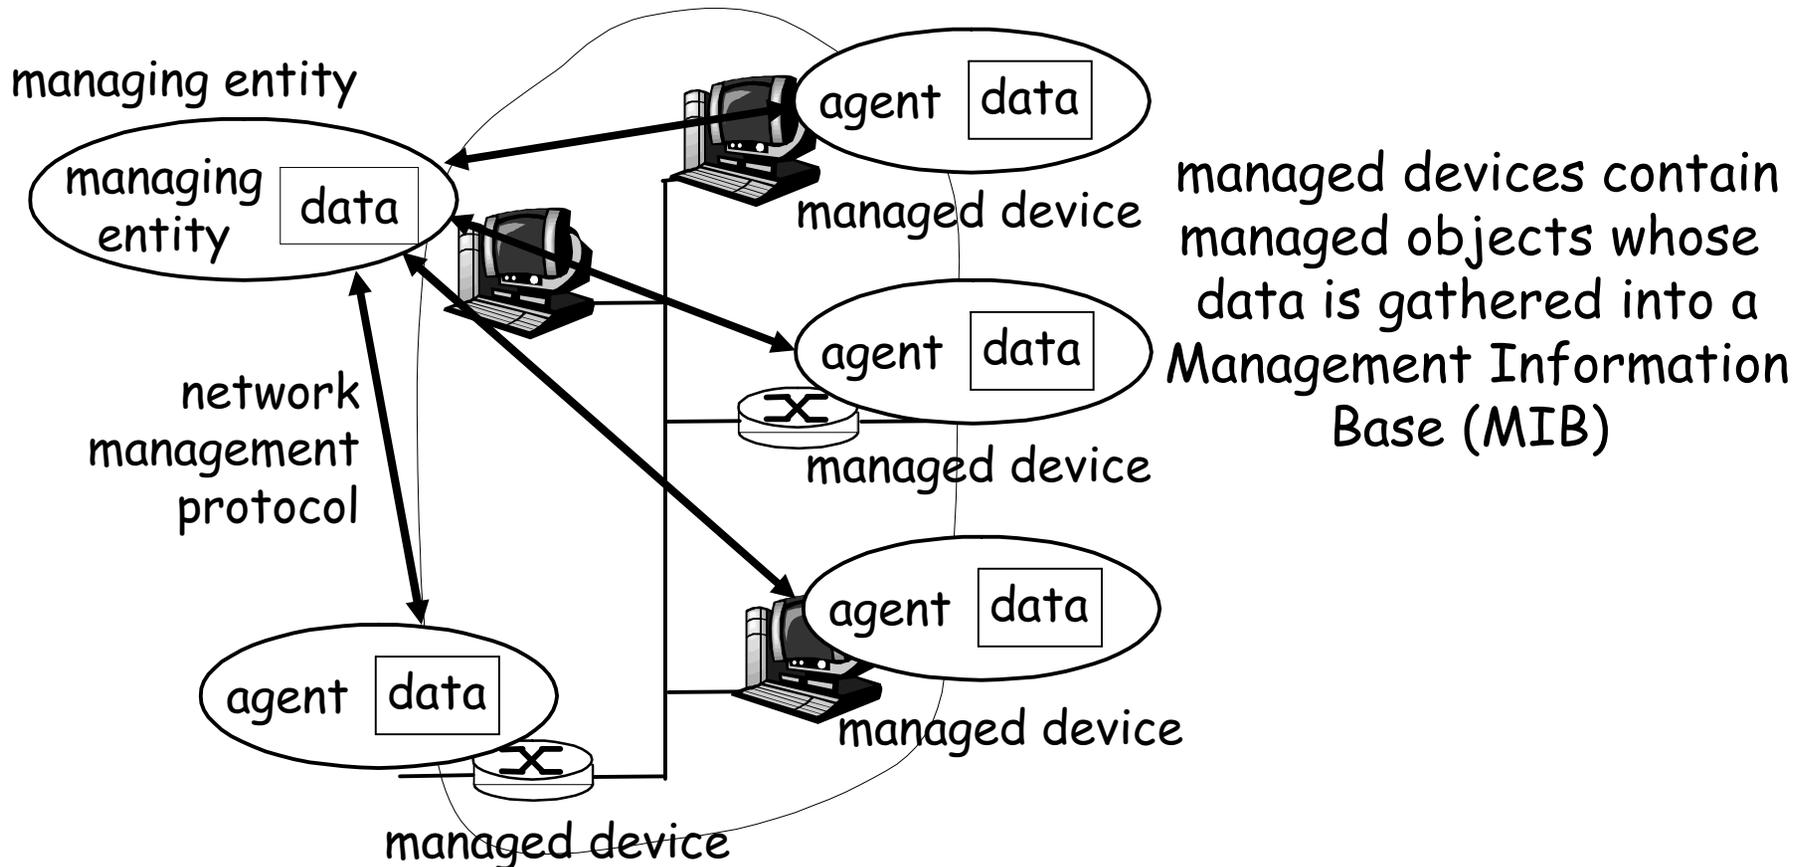
- ◆ autonomous systems (aka "network"): 100s or 1000s of interacting hardware/software components
- ◆ other complex systems requiring monitoring, control:
 - jet airplane
 - nuclear power plant
 - others?



"Network management includes the deployment, integration and coordination of the hardware, software, and human elements to monitor, test, poll, configure, analyze, evaluate, and control the network and element resources to meet the real-time, operational performance, and Quality of Service requirements at a reasonable cost."

Infrastructure for network management

definitions:



Network Management standards

SNMP: Simple Network Management Protocol

- ◆ Internet roots (SGMP)
- ◆ started simple
- ◆ deployed, adopted rapidly
- ◆ growth: size, complexity
- ◆ currently: SNMP V3
- ◆ *de facto* network management standard

SNMP overview: 4 key parts

- ◆ Management information base (MIB):
 - distributed information store of network management data
- ◆ Structure of Management Information (SMI):
 - data definition language for MIB objects
- ◆ SNMP protocol
 - convey manager<->managed object info, commands
- ◆ security, administration capabilities
 - major addition in SNMPv3

MIB example: UDP module

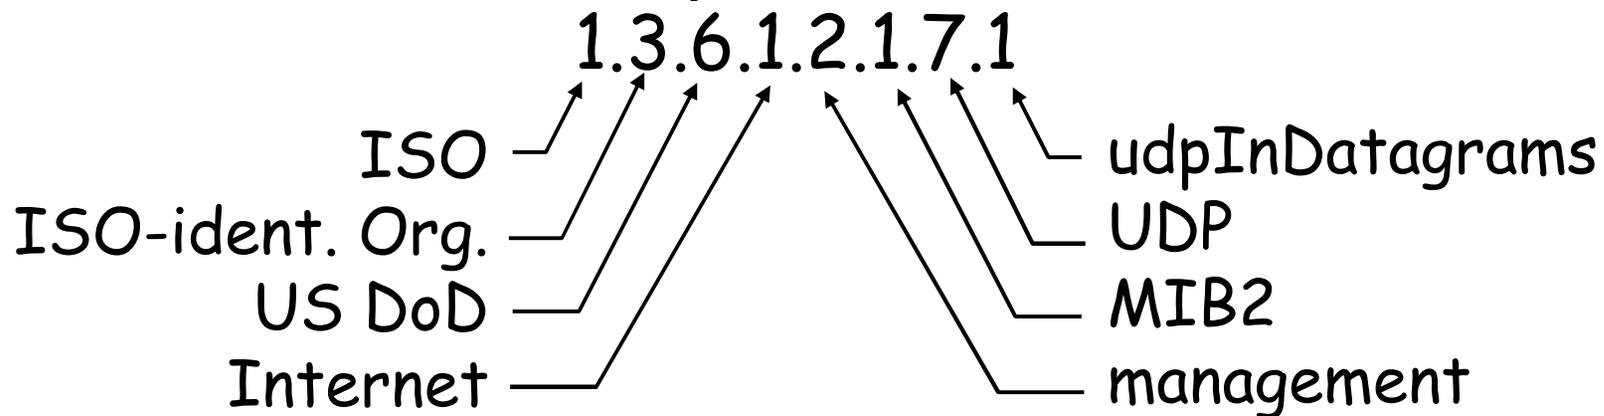
<u>Object ID</u>	<u>Name</u>	<u>Type</u>	<u>Comments</u>
1.3.6.1.2.1.7.1	UDPInDatagrams	Counter32	total # datagrams delivered at this node
1.3.6.1.2.1.7.2	UDPNoPorts	Counter32	# undeliverable datagrams no app at port
1.3.6.1.2.1.7.3	UDInErrors	Counter32	# undeliverable datagrams all other reasons
1.3.6.1.2.1.7.4	UDPOutDatagrams	Counter32	# datagrams sent
1.3.6.1.2.1.7.5	udpTable	SEQUENCE	one entry for each port in use by app, gives port # and IP address

SNMP Naming

question: how to name every possible standard object (protocol, data, more..) in every possible network standard??

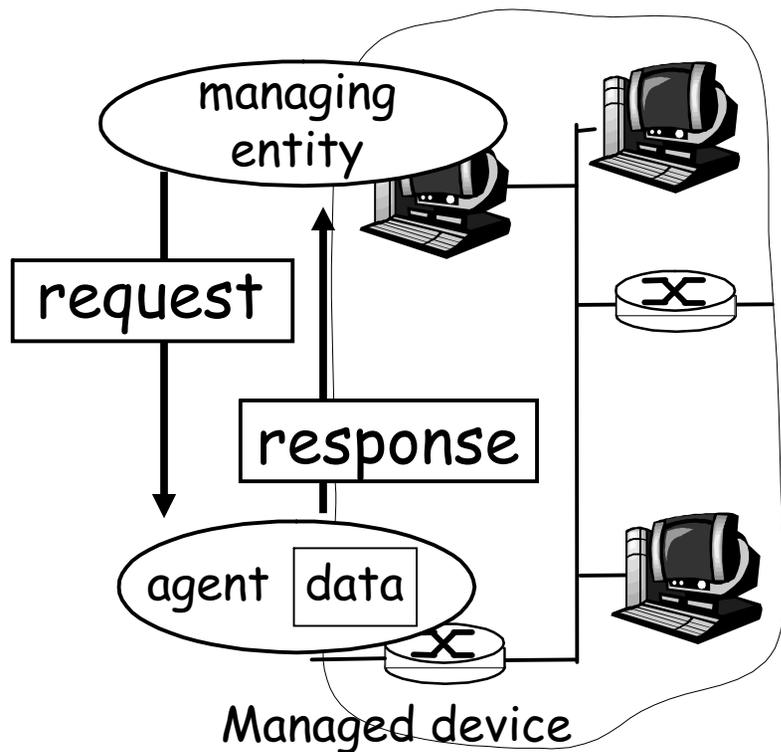
answer: *ISO Object Identifier tree:*

- hierarchical naming of all objects
- each branchpoint has name, number

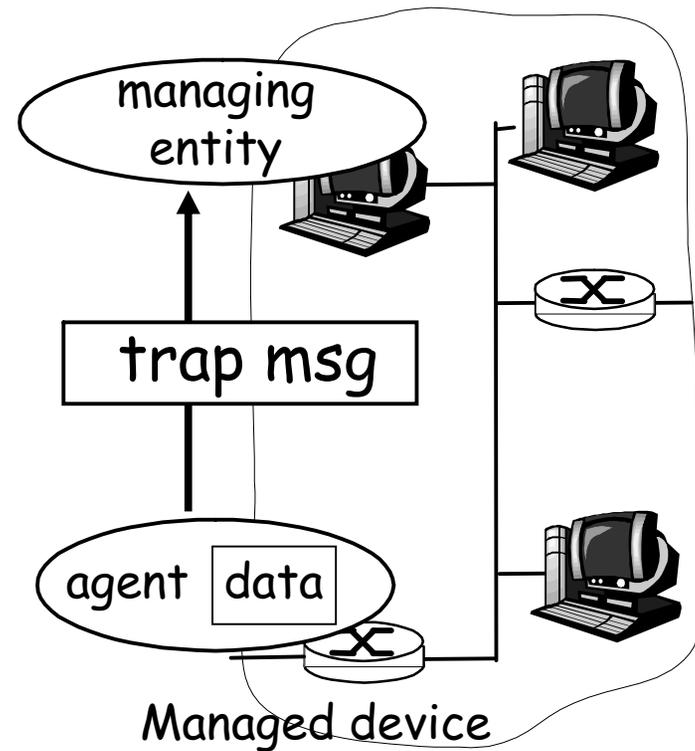


SNMP protocol

Two ways to convey MIB info, commands:



request/response mode



trap mode

SNMP protocol: message types

<u>Message type</u>	<u>Function</u>
GetRequest GetNextRequest GetBulkRequest	Mgr-to-agent: "get me data" (instance,next in list, block)
InformRequest	Mgr-to-Mgr: here's MIB value
SetRequest	Mgr-to-agent: set MIB value
Response	Agent-to-mgr: value, response to Request
Trap	Agent-to-mgr: inform manager of exceptional event

SNMP security and administration

- ◆ encryption: DES-encrypt SNMP message
- ◆ authentication: compute, send $\text{MIC}(m,k)$:
compute hash (MIC) over message (m),
secret shared key (k)
- ◆ protection against playback: use nonce
- ◆ view-based access control
 - SNMP entity maintains database of access rights, policies for various users
 - database itself accessible as managed object!