

UNIT 8. SECONDARY STORAGE, COMPUTER SECURITY

8.1 DISK STRUCTURE

Disks provide the bulk of secondary storage for modern computer systems. Magnetic tape was used as an early secondary-storage medium, but the access time is much slower than for disks. Modern disk drives are addressed as large one-dimensional arrays of **logical blocks**, where the logical block is the smallest unit of transfer. The size of a logical block is usually 512 bytes, although some disks can be low-level formatted to choose a different logical block size, such as 1,024 bytes.

The one-dimensional array of logical blocks is mapped onto the sectors of the disk sequentially. Sector 0 is the first sector of the first track on the outermost cylinder. The mapping proceeds in order through that track, then through the rest of the tracks in that cylinder, and then through the rest of the cylinders from outermost to innermost.

8.2 DISK SCHEDULING

One of the responsibilities of the OS is to use the hardware efficiently. This will help in achieving fast access time and disk bandwidth. The access time has two major components:

- The **seek time** is the time for the disk arm to move the heads to the cylinder containing the desired sector.
- The **rotational latency** is the additional time waiting for the disk to rotate the desired sector to the disk head.

The **disk bandwidth** is the total number of bytes transferred, divided by the total time between the first request for service and the completion of the last transfer. We can improve both the access time and the bandwidth by scheduling the servicing of disk I/O requests in a good order.

Whenever a process needs I/O to or from the disk, it issues a system call to the operating system. The request specifies several pieces of information:

- Whether this operation is input or output
- What the disk address for the transfer is
- What the memory address for the transfer is
- What the number of bytes to be transferred is

If the desired disk drive and controller are available, the request can be serviced immediately. If the drive or controller is busy, any new requests for service will be placed on the queue of pending requests for that drive. For a multiprogramming system with many processes, the disk queue may often have several pending requests. Thus, when one request is completed, the operating system chooses which pending request to service next.

8.2.1 FCFS Scheduling

The simplest form of disk scheduling is, of course, the first-come, first-served (FCFS) algorithm. This algorithm is intrinsically fair, but it generally does not provide the fastest service.

Example: Consider a disk queue with requests for I/O to blocks on cylinders:

98, 183, 37, 122, 14, 124, 65, 67

in that order. There are 200 cylinders numbered from 0 to 199. The disk head is initially at the cylinder 53. Compute total head movements.

Solution: Starting with cylinder 53, it will first move from 53 to 98, then to 183, 37, 122, 14, 124, 65, and finally to 67 as shown in Figure 8.1.

Head movement from 53 to 98 = 45
Head movement from 98 to 183 = 85
Head movement from 183 to 37 = 146
Head movement from 37 to 122 = 85
Head movement from 122 to 14 = 108
Head movement from 14 to 124 = 110
Head movement from 124 to 65 = 59
Head movement from 65 to 67 = 2

Total head movement = 640




Figure 8.1 FCFS Scheduling

8.2.2 SSTF Scheduling

The SSTF (shortest-seek-time-first) algorithm selects the request with the minimum seek time from the current head position. Since seek time increases with the number of cylinders traversed by the head, SSTF chooses the pending request closest to the current head position.

Example: Consider a disk queue with requests for I/O to blocks on cylinders:

98, 183, 37, 122, 14, 124, 65, 67

in that order. There are 200 cylinders numbered from 0 to 199. The disk head is initially at the cylinder 53. Compute total head movements.

Solution: The closest request to the initial head position 53 is at cylinder 65. Once we are at cylinder 65, the next closest request is at cylinder 67. From there, the request at cylinder 37 is closer than 98, so 37 is served next. Continuing, we service the request at cylinder 14, then 98, 122, 124, and finally 183. It is shown in Figure 8.2.

Head movement from 53 to 65 = 12
Head movement from 65 to 67 = 2
Head movement from 67 to 37 = 30
Head movement from 37 to 14 = 23
Head movement from 14 to 98 = 84
Head movement from 98 to 122 = 24
Head movement from 122 to 124 = 2
Head movement from 124 to 183 = 59
Total head movement = 236




Figure 8.2 SSTF Scheduling

8.2.3 SCAN Scheduling

In the SCAN algorithm, the disk arm starts at one end of the disk, and moves toward the other end, servicing requests as it reaches each cylinder, until it gets to the other end of the disk. At the other end, the direction of head movement is reversed, and servicing continues. The head continuously scans back and forth across the disk.

Example: Consider a disk queue with requests for I/O to blocks on cylinders:

98, 183, 37, 122, 14, 124, 65, 67

in that order. There are 200 cylinders numbered from 0 to 199. The disk head is initially at the cylinder 53. Compute total head movements.

Solution: Before applying SCAN algorithm, we need to know the current direction of head movement. Assume that disk arm is moving toward 0, the head will service 37 and then 14.

At cylinder 0, the arm will reverse and will move toward the other end of the disk, servicing the requests at 65,67,98, 122, 124, and 183. It is shown in Figure 8.3.

Head movement from 53 to 37 = 16

Head movement from 37 to 14 = 23

Head movement from 14 to 0 = 14

Head movement from 0 to 65 = 65

Head movement from 65 to 67 = 2

Head movement from 67 to 98 = 31

Head movement from 98 to 122 = 24

Head movement from 122 to 124 = 2

Head movement from 124 to 183 = 59

Total head movement = 236




Figure 8.3 SCAN disk scheduling

8.2.4 C-SCAN Scheduling

Circular SCAN (C-SCAN) scheduling is a variant of SCAN designed to provide a more uniform wait time. Like SCAN, C-SCAN moves the head from one end of the disk to the other, servicing requests along the way. When the head reaches the other end, however, it immediately returns to the beginning of the disk, without servicing any requests on the return trip (Figure 8.4). The C-SCAN scheduling algorithm essentially treats the cylinders as a circular list that wraps around from the final cylinder to the first one.

Example: Consider a disk queue with requests for I/O to blocks on cylinders:

98, 183, 37, 122, 14, 124, 65, 67

in that order. There are 200 cylinders numbered from 0 to 199. The disk head is initially at the cylinder 53. Compute total head movements.

Solution: Before applying C - SCAN algorithm, we need to know the current direction of head movement. Assume that disk arm is moving toward 199, the head will service 65, 67, 98, 122, 124, 183. Then it will move to 199 and the arm will reverse and move towards 0. While moving towards 0, it will not serve. But, after reaching 0, it will reverse again and then serve 14 and 37. It is shown in Figure 8.4.

Head movement from 53 to 65 = 12

Head movement from 65 to 67 = 2

Head movement from 67 to 98 = 31

Head movement from 98 to 122 = 24

Head movement from 122 to 124 = 2

Head movement from 124 to 183 = 59

Head movement from 183 to 199 = 16

Head movement from 199 to 0 = 199

Head movement from 0 to 14 = 14

Head movement from 14 to 37 = 23

Total head movement = 382




Figure 8.4 C-SCAN Disk Scheduling

8.2.5 Look Scheduling

As we described them, both SCAN and C-SCAN move the disk arm across the full width of the disk. In practice, neither algorithm is implemented this way. More commonly, the arm goes only as far as the final request in each direction. Then, it reverses direction immediately, without going all the way to the end of the disk. These versions of SCAN and C-SCAN are called LOOK and C-LOOK scheduling, because they look for a request before continuing to move in a given direction.

Example: Consider a disk queue with requests for I/O to blocks on cylinders:

98, 183, 37, 122, 14, 124, 65, 67

in that order. There are 200 cylinders numbered from 0 to 199. The disk head is initially at the cylinder 53. Compute total head movements.

Solution: Assume that disk arm is moving toward 199, the head will service 65, 67, 98, 122, 124, 183. Then the arm will reverse and move towards 14. Then it will serve 37. It is shown in Figure 8.5.

Head movement from 53 to 65 = 12
Head movement from 65 to 67 = 2
Head movement from 67 to 98 = 31
Head movement from 98 to 122 = 24
Head movement from 122 to 124 = 2
Head movement from 124 to 183 = 59
Head movement from 183 to 14 = 169
Head movement from 14 to 37 = 23

Total head movement = 322




Figure 8.5 Look Disk Scheduling

8.3 DISK MANAGEMENT

The OS is responsible for several other aspects of disk management as well.

8.3.1 Disk Formatting

A new magnetic disk is a blank slate: It is just platters of a magnetic recording material. Before a disk can store data, it must be divided into sectors that the disk controller can read and write. This process is called low-level formatting (or physical formatting). Low-level formatting fills the disk with a special data structure for each sector.

To use a disk to hold files, the OS needs to record its own data structures on the disk. It does so in two steps. The first step is to **partition** the disk into one or more groups of cylinders. The OS treats each partition as though it were a separate disk. After partitioning, the second step is **logical formatting** (or creation of a file system). In this step, the OS

stores the initial file-system data structures onto the disk. These data structures may include maps of free and allocated space and an initial empty directory.

8.3.2 Boot Blocks

For a computer to start running, it needs to have an initial program to run. This bootstrap program initializes all aspects of the system, from CPU registers to device controllers and the contents of main memory, and then starts the OS. To do its job, the bootstrap program finds the OS kernel on disk, loads that kernel into memory, and jumps to an initial address to begin the operating-system execution.

8.3.3 Bad Blocks

Because disks have moving parts and small tolerances, they are prone to failure. Sometimes the failure is complete, and the disk needs to be replaced, and its contents restored from backup media to the new disk. More frequently, one or more sectors become defective. Most disks even come from the factory with bad blocks. Depending on the disk and controller in use, these blocks are handled in a variety of ways:

- On simple disks with IDE, bad blocks are handled manually (like formatting the disk)
- The controller maintains a list of bad blocks on the disk. The list is initialized during the low-level format at the factory, and is updated over the life of the disk.

8.4 THE SECURITY PROBLEM

We have discussed earlier that each file/directory can be secured by controlling access to them. But, total security cannot be achieved by this. Security violations of the system can be categorized as intentional (malicious) or accidental. It is easier to protect against accidental misuse. Malicious access has following forms:

- Unauthorized reading of data
- Unauthorized modification of data
- Unauthorized destruction of data
- Preventing legitimate use of the system

To protect the system, we must take security measures at four levels:

1. **Physical:** The site or sites containing the computer systems must be physically secured against armed or surreptitious entry by intruders.
2. **Human:** Users must be screened carefully to reduce the chance of authorizing a user who then gives access to an intruder.
3. **Network:** Data being transferred via network is prone to attack.
4. **Operating System:** The system must protect itself security breaches.

8.5 USER AUTHENTICATION

The major security problem for OS is authentication. The protection system depends on an ability identify the programs and processes currently executing. Generally authentication is based on one or more of three items:

- User possession (a key or a card)
- User knowledge (identifier or password)
- User attribute (fingerprint, retina, signature etc)

8.5.1 Passwords

Passwords are used to protect the data in the computer, when there are no complete protection schemes. User is asked to provide the username and password during the access. If the password matches with the one stored in the system, access is allowed.

8.5.2 Password Vulnerabilities

Though passwords are popularly used because of its easy usage and understanding, it is vulnerable to threats. They can be guessed, accidentally exposed, illegally transferred from one to other etc. To avoid these various possibilities can be tried:

- keeping non-guessable passwords
- keep changing the passwords frequently
- keeping encrypted passwords
- using One – Time Passwords (OTP)
- using biometrics authentication

8.6 PROGRAM THREATS

When a program written by one user is used by another user, misuse and unexpected behavior may occur. Following are the common methods by which such behavior may occur:

- **Trojan Horse:** A code segment that misuses its environment is called as Trojan horse. For example, a hacker would have created a fake login screen on your computer. When you enter username and password, it will store the data, but displays the message as 'Login failed'. You may feel that you have typed wrongly. That code will now terminate and actual login pops-up. You will once again give login credentials and the system logs-in. Such a code which created fake screen is Trojan Horse.
- **Trap Door:** The designer of a program or system may leave a loop-hole in the software purposefully. Only he/she knows it and tries to collect information from a particular user without their knowledge. This type of security breach is known as trap door. A trap door can also be included in a compiler.
- **Stack and Buffer Overflow:** A bug is introduced in a program such a way that, it tries to access stack or buffer beyond its limits. Hence, memory safety is violated.

8.7 SYSTEM THREATS

Most OS allows processes to spawn (creating child process) other processes. In such a situation, it is possible that OS resources and user files are misused. The methods for misusing are discussed here:

- **Worms:** A worm is a process that uses the spawn mechanism to attack system performance. It creates multiple copies of itself using system resources and locking all other processes from using the system.
- **Viruses:** It is a code embedded within a legitimate program. They spread into other programs and mess-up the system – modifying/destroying files, causing system crash etc.
- **Denial of Service:** Here, the hacker does not gain any information, but disables legitimate user from accessing system.

CASE STUDY OF LINUX OS

8.8 LINUX HISTORY

Linux is a modern, free operating system based on UNIX standards. It is developed as a small but self-contained kernel in 1991 by Linus Torvalds, with the major design goal of UNIX compatibility. Its history has been one of collaboration by many users from all around the world, corresponding almost exclusively over the Internet. It has been designed to run efficiently and reliably on common PC hardware, but also runs on a variety of other platforms. The core Linux operating system kernel is entirely original, but it can run much existing free UNIX software, resulting in an entire UNIX-compatible operating system free from proprietary code.

8.9 DESIGN PRINCIPLES

Linux is a multiuser, multitasking system with a full set of UNIX-compatible tools. Its file system adheres to traditional UNIX semantics, and it fully implements the standard UNIX networking model. Main design goals are speed, efficiency, and standardization. Linux is designed to be compliant with the relevant POSIX documents; at least two Linux distributions have achieved official POSIX certification. The Linux programming interface adheres to the SVR4 UNIX semantics, rather than to BSD behavior.

8.9.1 Components of Linux System

There are three major components in Linux:

- **Kernel:** It is responsible for maintaining all the important abstractions of the OS, including virtual memory and processes. Kernel code executes in *kernel mode* with full access to all the physical resources of the computer. All kernel code and data structures are kept in the same single address space.
- **System Libraries:** They define a standard set of functions through which applications can interact with the kernel. They also implement much of the OS functionality that does not need the full privileges of kernel code.
- **System Utilities:** These are the programs which perform individual specialized management tasks.

Figure 8.6 illustrates various components of Linux system.

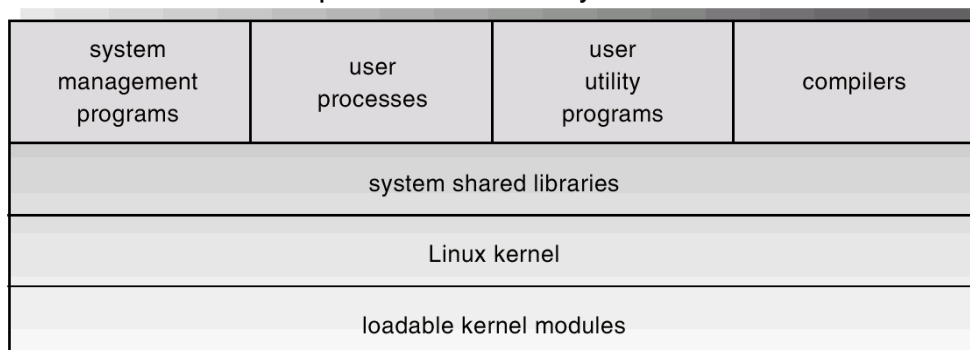


Figure 8.6 Components of Linux System

8.10 KERNEL MODULES

Kernel modules are the sections of kernel code that can be compiled, loaded, and unloaded independent of the rest of the kernel. A kernel module may typically implement a device driver, a file system, or a networking protocol. The module interface allows third parties to write and distribute, on their own terms, device drivers or file systems that could not be distributed under the GPL (General Public Library). Kernel modules allow a Linux system to be set up with a standard, minimal kernel, without any extra device drivers built in. Three components to Linux module support are discussed here.

- **Module Management:** It supports to load modules into memory and letting them communicate to the rest of the kernel. Module loading is split into two separate sections:
 - Managing sections of module code in kernel memory
 - Handling symbols that modules are allowed to reference

The module requestor manages loading requested, but currently unloaded, modules; it also regularly queries the kernel to see whether a dynamically loaded module is still in use, and will unload it when it is no longer actively needed.

- **Driver Registration:** Allows modules to tell the rest of the kernel that a new driver has become available. The kernel maintains dynamic tables of all known drivers, and provides a set of routines to allow drivers to be added to or removed from these tables at any time. Registration tables include the following items:
 - Device drivers
 - File systems
 - Network protocols
 - Binary format
- **Conflict Resolution:** A mechanism that allows different device drivers to reserve hardware resources and to protect those resources from accidental use by another driver. The conflict resolution module aims to:
 - Prevent modules from clashing over access to hardware resources
 - Prevent *autoprobes* from interfering with existing device drivers
 - Resolve conflicts with multiple drivers trying to access the same hardware

8.11 PROCESS MANAGEMENT

A process is the basic context within which all user-requested activity is serviced within the OS. To be compatible with other UNIX systems, Linux must use a process model similar to those of other versions of UNIX.

8.11.1 The Fork/Exec Process Model

UNIX process management separates the creation of processes and the running of a new program into two distinct operations. The **fork** system call creates a new process. A new program is run after a call to **exec**. Under UNIX, a process encompasses all the information that the operating system must maintain to track the context of a single execution of a single program.

Under Linux, process properties fall into three groups:

- **Process Identity:** The process identity consists of mainly following items:

- **Process ID (PID):** The unique identifier for the process; used to specify processes to the OS when an application makes a system call to signal, modify, or wait for another process.
- **Credentials:** Each process must have an associated user ID and one or more group IDs that determine the process's rights to access system resources and files.
- **Personality:** Not traditionally found on UNIX systems, but under Linux each process has an associated personality identifier that can slightly modify the semantics of certain system calls. Used primarily by emulation libraries to request that system calls be compatible with certain specific flavors of UNIX.
- **Process environment:** The process's environment is inherited from its parent, and is composed of two null-terminated vectors:
 - The **argument vector** lists the command-line arguments used to invoke the running program; conventionally starts with the name of the program itself
 - The **environment vector** is a list of "NAME=VALUE" pairs that associates named environment variables with arbitrary textual values.

Passing environment variables among processes and inheriting variables by a process's children are flexible means of passing information to components of the user-mode system software. The environment-variable mechanism provides a customization of the operating system that can be set on a per-process basis, rather than being configured for the system as a whole.

- **Context:** It is the (constantly changing) state of a running program at any point in time. There are various parts:
 - **Scheduling context** is the most important part of the process context; it is the information that the scheduler needs to suspend and restart the process.
 - **Accounting:** The kernel maintains accounting information about the resources currently being consumed by each process, and the total resources consumed by the process in its lifetime so far.
 - **File table** is an array of pointers to kernel file structures. When making file I/O system calls, processes refer to files by their index into this table.
 - **File System Context:** Whereas the file table lists the existing open files, the file-system context applies to requests to open new files. The current root and default directories to be used for new file searches are stored here.
 - **Signal-handler table** defines the routine in the process's address space to be called when specific signals arrive.
 - **Virtual-memory context** of a process describes the full contents of its private address space.

8.11.2 Processes and Threads

Linux uses the same internal representation for processes and threads; a thread is simply a new process that happens to share the same address space as its parent. A distinction is only made when a new thread is created by the **clone** system call:

- **fork** creates a new process with its own entirely new process context

- **clone** creates a new process with its own identity, but that is allowed to share the data structures of its parent

Using **clone** gives an application fine-grained control over exactly what is shared between two threads.

8.12 SCHEDULING

The job of allocating CPU time to different tasks within an OS. While scheduling is normally thought of as the running and interrupting of processes, in Linux, scheduling also includes the running of the various kernel tasks. Running kernel tasks encompasses both tasks that are requested by a running process and tasks that execute internally on behalf of a device driver.

Various aspects of scheduling in Linux are discussed here.

- **Kernel Synchronization:** A request for kernel-mode execution can occur in two ways:
 - A running program may request an operating system service, either explicitly via a system call, or implicitly, for example, when a page fault occurs.
 - A device driver may deliver a hardware interrupt that causes the CPU to start executing a kernel-defined handler for that interrupt.

Kernel synchronization requires a framework that will allow the kernel's critical sections to run without interruption by another critical section.

Linux uses two techniques to protect critical sections:

- Normal kernel code is non-preemptible : when a time interrupt is received while a process is executing a kernel system service routine, the kernel's **need_resched** flag is set so that the scheduler will run once the system all has completed and control is about to be returned to user mode.
 - The second technique applies to critical sections that occur in an interrupt service routines. By using the processor's interrupt control hardware to disable interrupts during a critical section, the kernel guarantees that it can proceed without the risk of concurrent access of shared data structures.
- **Process Scheduling:** Linux uses two process-scheduling algorithms:
 - A time-sharing algorithm for fair preemptive scheduling between multiple processes
 - A real-time algorithm for tasks where absolute priorities are more important than fairness

A process's scheduling class defines which algorithm to apply. For time-sharing processes, Linux uses a prioritized, credit based algorithm.

- **Symmetric Multiprocessing:** Linux 2.0 was the first Linux kernel to support SMP hardware; separate processes or threads can execute in parallel on separate processors. To preserve the kernel's nonpreemptible synchronization requirements, SMP imposes the restriction, via a single kernel spinlock, that only one processor at a time may execute kernel-mode code.

8.13 MEMORY MANAGEMENT

Linux's physical memory-management system deals with allocating and freeing pages, groups of pages, and small blocks of memory. It has additional mechanisms for handling virtual memory, memory mapped into the address space of running processes.

- **Management of Physical Memory:** The page allocator allocates and frees all physical pages; it can allocate ranges of physically-contiguous pages on request. The allocator uses a *buddy-heap* algorithm to keep track of available physical pages:
 - Each allocatable memory region is paired with an adjacent partner.
 - Whenever two allocated partner regions are both freed up they are combined to form a larger region.
 - If a small memory request cannot be satisfied by allocating an existing small free region, then a larger free region will be subdivided into two partners to satisfy the request.

Memory allocations in the Linux kernel occur either statically (drivers reserve a contiguous area of memory during system boot time) or dynamically (via the page allocator).

- **Management of Virtual Memory:** The VM system maintains the address space visible to each process: It creates pages of virtual memory on demand, and manages the loading of those pages from disk or their swapping back out to disk as required. The VM manager maintains two separate views of a process's address space:
 - A logical view describing instructions concerning the layout of the address space. The address space consists of a set of nonoverlapping regions, each representing a continuous, page-aligned subset of the address space.
 - A physical view of each address space which is stored in the hardware page tables for the process.

Virtual memory regions are characterized by:

- The backing store, which describes from where the pages for a region come; regions are usually backed by a file or by nothing (*demand-zero* memory)
- The region's reaction to writes (page sharing or copy-on-write).

The kernel creates a new virtual address space

1. When a process runs a new program with the **exec** system call
2. Upon creation of a new process by the **fork** system call

- **Executing and Loading User Programs:** Linux maintains a table of functions for loading programs; it gives each function the opportunity to try loading the given file when an exec system call is made. The registration of multiple loader routines allows Linux to support both the ELF and a.out binary formats. Initially, binary-file pages are mapped into virtual memory; only when a program tries to access a given page will a page fault result in that page being loaded into physical memory. An ELF-format binary file consists of a header followed by several page-aligned sections; the ELF loader works by reading the header and mapping the sections of the file into separate regions of virtual memory.

8.14 FILE SYSTEMS

To the user, Linux's file system appears as a hierarchical directory tree obeying UNIX semantics. Internally, the kernel hides implementation details and manages the multiple

different file systems via an abstraction layer, that is, the virtual file system (VFS). The Linux VFS is designed around object-oriented principles and is composed of two components:

- A set of definitions that define what a file object is allowed to look like
 - The inode-object and the file-object structures represent individual files
 - the file system object represents an entire file system
- A layer of software to manipulate those objects.

8.15 INPUT AND OUTPUT

The Linux device-oriented file system accesses disk storage through two caches:

- Data is cached in the page cache, which is unified with the virtual memory system
- Metadata is cached in the buffer cache, a separate cache indexed by the physical disk block.

Linux splits all devices into three classes:

- block devices allow random access to completely independent, fixed size blocks of data
- character devices include most other devices; they don't need to support the functionality of regular files.
- network devices are interfaced via the kernel's networking subsystem

8.16 INTER-PROCESS COMMUNICATION

Like UNIX, Linux informs processes that an event has occurred via signals. There is a limited number of signals, and they cannot carry information: Only the fact that a signal occurred is available to a process. The Linux kernel does not use signals to communicate with processes with are running in kernel mode, rather, communication within the kernel is accomplished via scheduling states and wait-queue structures.

Passing of Data among Processes: The pipe mechanism allows a child process to inherit a communication channel to its parent; data written to one end of the pipe can be read by the other. Shared memory offers an extremely fast way of communicating; any data written by one process to a shared memory region can be read immediately by any other process that has mapped that region into its address space. To obtain synchronization, however, shared memory must be used in conjunction with another Interprocess-communication mechanism.
