SUBJECT-Operating system

INTRODUCTION [UNIT-I]

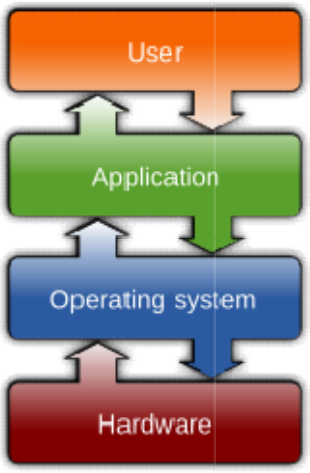

-
-

An **operating system** (**OS**) is system software that manages computer hardware and software resources, and provides common services for computer programs.

Time-sharing operating systems schedule tasks for efficient use of the system and may also include accounting software for cost allocation of processor time, mass storage, peripherals, and other resources.

For hardware functions such as input and output and memory allocation, the operating system acts as an intermediary between programs and the computer hardware,[1][2] although the application code is usually executed directly by the hardware and frequently makes system calls to an OS function or is interrupted by it. Operating systems are found on many devices that contain a computer – from cellular phones and video game consoles to web servers and supercomputers.

In the personal computer market, as of September 2024, Microsoft Windows holds a dominant market share of around 73%. macOS by Apple Inc. is in second place (15%), Linux is in third place (5%), and ChromeOS is in fourth place (2%).[3] In the mobile sector (including smartphones and tablets), as of September 2023, Android's share is 68.92%, followed by Apple's iOS and iPadOS with 30.42%, and other operating systems with .66%.[4] Linux distributions are dominant in the server and supercomputing sectors. Other specialized classes of operating systems (special-purpose operating systems),[5][6] such as embedded and real-time systems, exist

for many applications. Security-focused operating systems also exist. Some operating systems have low system requirements (e.g. light-weight Linux distribution). Others may have higher system requirements.

Some operating systems require installation or may come pre-installed with purchased computers (OEM-installation), whereas others may run directly from media (i.e. live CD) or flash memory (i.e. USB stick).

Definition and purpose

An operating system is difficult to define,[7] but has been called "the layer of software that manages a computer's resources for its users and their applications".[8] Operating systems include the software that is always running, called a kernel—but can include other software as well.[7][9] The two other types of programs that can run on a computer are system programs—which are associated with the operating system, but may not be part of the kernel—and applications—all other software.[9]

There are three main purposes that an operating system fulfills:[10]

- Operating systems allocate resources between different applications, deciding when they will receive central processing unit (CPU) time or space in memory.[10] On modern personal computers, users often want to run several applications at once. In order to ensure that one program cannot monopolize the computer's limited hardware resources, the operating system gives each application a share of the resource, either in time (CPU) or space (memory).[11][12] The operating system also must isolate applications from each other to protect them from errors and security vulnerability in another application's code, but enable communications between different applications.[13]
- Operating systems provide an interface that abstracts the details of accessing hardware details (such as physical memory) to make things easier for programmers.[10][14] Virtualization also enables the operating system to mask limited hardware resources; for example, virtual memory can provide a program with the illusion of nearly unlimited memory that exceeds the computer's actual memory.[15]
- Operating systems provide common services, such as an interface for accessing network and disk devices. This enables an application to be run on different hardware without needing to be rewritten.[16] Which services to include in an operating system varies greatly, and this functionality makes up the great majority of code for most operating systems.[17]

Types of operating systems

Multicomputer operating systems

With multiprocessors multiple CPUs share memory. A multicomputer or cluster computer has multiple CPUs, each of which has its own memory. Multicomputers were developed because large multiprocessors are difficult to engineer and prohibitively expensive;[18] they are universal in cloud computing because of the size of the machine needed.[19] The different CPUs often need to send and receive messages to each

other;[20] to ensure good performance, the operating systems for these machines need to minimize this copying of packets.[21] Newer systems are often multiqueue—separating groups of users into separate queues—to reduce the need for packet copying and support more concurrent users.[22] Another technique is remote direct memory access, which enables each CPU to access memory belonging to other CPUs.[20] Multicomputer operating systems often support remote procedure calls where a CPU can call a procedure on another CPU,[23] or distributed shared memory, in which the operating system uses virtualization to generate shared memory that does not actually exist.[24]

Distributed systems

A distributed system is a group of distinct, networked computers—each of which might have their own operating system and file system. Unlike multicomputers, they may be dispersed anywhere in the world.[25] Middleware, an additional software layer between the operating system and applications, is often used to improve consistency. Although it functions similarly to an operating system, it is not a true operating system.[26]

Embedded

Embedded operating systems are designed to be used in embedded computer systems, whether they are internet of things objects or not connected to a network. Embedded systems include many household appliances. The distinguishing factor is that they do not load user-installed software. Consequently, they do not need protection between different applications, enabling simpler designs. Very small operating systems might run in less than 10 kilobytes,[27] and the smallest are for smart cards.[28] Examples include Embedded Linux, QNX, VxWorks, and the extra-small systems RIOT and TinyOS.[29]

Real-time

A real-time operating system is an operating system that guarantees to process events or data by or at a specific moment in time. Hard real-time systems require exact timing and are common in manufacturing, avionics, military, and other similar uses.[29] With soft real-time systems, the occasional missed event is acceptable; this category often includes audio or multimedia systems, as well as smartphones.[29] In order for hard real-time systems be sufficiently exact in their timing, often they are just a library with no protection between applications, such as eCos.[29]

Hypervisor

A hypervisor is an operating system that runs a virtual machine. The virtual machine is unaware that it is an application and operates as if it had its own hardware.[15][30] Virtual machines can be paused, saved, and resumed, making them useful for operating systems research, development,[31] and debugging.[32] They also enhance portability by enabling applications to be run on a computer even if they are not compatible with the base operating system.[15]

Library

A *library operating system* (libOS) is one in which the services that a typical operating system provides, such as networking, are provided in the form of libraries and composed with a single application and configuration code to construct a unikernel: [33] a specialized (only the absolute necessary pieces of code are extracted from libraries and bound together [34]), single address space, machine image that can be deployed to cloud or embedded environments.

The operating system code and application code are not executed in separated protection domains (there is only a single application running, at least conceptually, so there is no need to prevent interference between applications) and OS services are accessed via simple library calls (potentially inlining them based on compiler thresholds), without the usual overhead of context switches, [35] in a way similarly to embedded and real-time OSes. Note that this overhead is not negligible: to the direct cost of mode switching it's necessary to add the indirect pollution of important processor structures (like CPU caches, the instruction pipeline, and so on) which affects both user-mode and kernel-mode performance. [36]

History

IBM System/360 Model 50 operator's console and CPU; the operator's console is a terminal used by the operating system to communicate with the operator.

The first computers in the late 1940s and 1950s were directly programmed either with plugboards or with machine code inputted on media such as punch cards, without programming languages or operating systems.[37] After the introduction of the transistor in the mid-1950s, mainframes began to be built. These still needed professional operators[37] who manually do what a modern operating system would do, such as scheduling programs to run,[38] but mainframes still had rudimentary operating systems such as Fortran Monitor System (FMS) and IBSYS.[39] In the 1960s, IBM introduced the first series of intercompatible computers (System/360). All of them ran the same operating system—OS/360—which consisted of millions of lines

of assembly language that had thousands of bugs. The OS/360 also was the first popular operating system to support multiprogramming, such that the CPU could be put to use on one job while another was waiting on input/output (I/O). Holding multiple jobs in memory necessitated memory partitioning and safeguards against one job accessing the memory allocated to a different one.[40]

Around the same time, teleprinters began to be used as terminals so multiple users could access the computer simultaneously. The operating system MULTICS was intended to allow hundreds of users to access a large computer. Despite its limited adoption, it can be considered the precursor to cloud computing. The UNIX operating system originated as a development of MULTICS for a single user.[41] Because UNIX's source code was available, it became the basis of other, incompatible operating systems, of which the most successful were AT&T's System V and the University of California's Berkeley Software Distribution (BSD).[42] To increase compatibility, the IEEE released the POSIX standard for operating system application programming interfaces (APIs), which is supported by most UNIX systems. MINIX was a stripped-down version of UNIX, developed in 1987 for educational uses, that inspired the commercially available, free software Linux. Since 2008, MINIX is used in controllers of most Intel microchips, while Linux is widespread in data centers and Android smartphones.[43]

Microcomputers

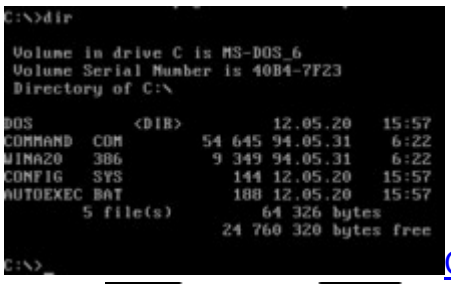

Command-line interface of the MS-DOS operating system

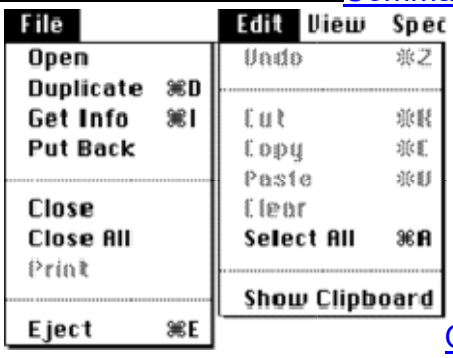

Graphical user interface of a Macintosh

The invention of large scale integration enabled the production of personal computers (initially called microcomputers) from around 1980.[44] For around five years, the CP/M (Control Program for Microcomputers) was the most popular operating system for microcomputers.[45] Later, IBM bought the DOS (Disk Operating System) from Microsoft. After modifications requested by IBM, the resulting system was called MS-DOS (MicroSoft Disk Operating System) and was widely used on IBM

microcomputers. Later versions increased their sophistication, in part by borrowing features from UNIX.[45]

Apple's Macintosh was the first popular computer to use a graphical user interface (GUI). The GUI proved much more user friendly than the text-only command-line interface earlier operating systems had used. Following the success of Macintosh, MS-DOS was updated with a GUI overlay called Windows. Windows later was rewritten as a stand-alone operating system, borrowing so many features from another (VAX VMS) that a large legal settlement was paid.[46] In the twenty-first century, Windows continues to be popular on personal computers but has less market share of servers. UNIX operating systems, especially Linux, are the most popular on enterprise systems and servers but are also used on mobile devices and many other computer systems.[47]

On mobile devices, Symbian OS was dominant at first, being usurped by BlackBerry OS (introduced 2002) and iOS for iPhones (from 2007). Later on, the open-source Android operating system (introduced 2008), with a Linux kernel and a C library (Bionic) partially based on BSD code, became most popular.[48]

Components

The components of an operating system are designed to ensure that various parts of a computer function cohesively. All user software must interact with the operating system to access hardware.

Kernel

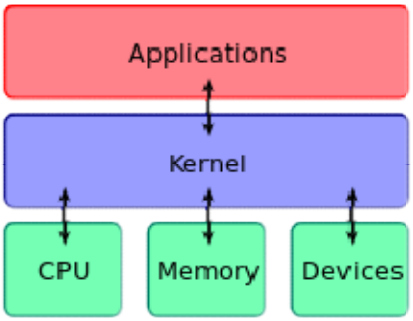

A kernel connects the application software to the hardware of a computer.

The kernel is the part of the operating system that provides protection between different applications and users. This protection is key to improving reliability by keeping errors isolated to one program, as well as security by limiting the power of malicious software and protecting private data, and ensuring that one program cannot monopolize the computer's resources.[49] Most operating systems have two modes of operation:[50] in user mode, the hardware checks that the software is only executing legal instructions, whereas the kernel has unrestricted powers and is not subject to these checks.[51] The kernel also manages memory for other processes and controls access to input/output devices.[52]

Program execution

The operating system provides an interface between an application program and the computer hardware, so that an application program can interact with the hardware only by obeying rules and procedures programmed into the operating system. The operating system is also a set of services which simplify development and execution of application programs. Executing an application program typically involves the creation of a process by the operating system kernel, which assigns memory space and other resources, establishes a priority for the process in multi-tasking systems, loads program binary code into memory, and initiates execution of the application program, which then interacts with the user and with hardware devices. However, in some systems an application can request that the operating system execute another application within the same process, either as a subroutine or in a separate thread, e.g., the **LINK** and **ATTACH** facilities of OS/360 and successors.

Interrupts

An interrupt (also known as an abort, exception, *fault*, signal,[53] or *trap*)[54] provides an efficient way for most operating systems to react to the environment. Interrupts cause the central processing unit (CPU) to have a control flow change away from the currently running program to an interrupt handler, also known as an interrupt service routine (ISR).[55][56] An interrupt service routine may cause the central processing unit (CPU) to have a context switch.[57][a] The details of how a computer processes an interrupt vary from architecture to architecture, and the details of how interrupt service routines behave vary from operating system to operating system.[58] However, several interrupt functions are common.[58] The architecture and operating system must:[58]

1. transfer control to an interrupt service routine.
2. save the state of the currently running process.
3. restore the state after the interrupt is serviced.

Software interrupt

A software interrupt is a message to a process that an event has occurred.[53] This contrasts with a *hardware interrupt* — which is a message to the central processing unit (CPU) that an event has occurred.[59] Software interrupts are similar to hardware interrupts — there is a change away from the currently running process.[60] Similarly, both hardware and software interrupts execute an interrupt service routine.

Software interrupts may be normally occurring events. It is expected that a time slice will occur, so the kernel will have to perform a context switch.[61] A computer program may set a timer to go off after a few seconds in case too much data causes an algorithm to take too long.[62]

Software interrupts may be error conditions, such as a malformed machine instruction.[62] However, the most common error conditions are division by zero and accessing an invalid memory address.[62]

Users can send messages to the kernel to modify the behavior of a currently running process.[62] For example, in the command-line environment, pressing the *interrupt character* (usually Control-C) might terminate the currently running process.[62]

To generate *software interrupts* for x86 CPUs, the INT assembly language instruction is available.[63] The syntax is `INT X`, where `X` is the offset number (in hexadecimal format) to the interrupt vector table.

Signal

To generate *software interrupts* in Unix-like operating systems, the `kill(pid,signum)` system call will send a signal to another process.[64] `pid` is the process identifier of the receiving process. `signum` is the signal number (in mnemonic format)[b] to be sent. (The abrasive name of `kill` was chosen because early implementations only terminated the process.)[65]

In Unix-like operating systems, *signals* inform processes of the occurrence of asynchronous events.[64] To communicate asynchronously, interrupts are required.[66] One reason a process needs to asynchronously communicate to another process solves a variation of the classic reader/writer problem.[67] The writer receives a pipe from the shell for its output to be sent to the reader's input stream.[68] The command-line syntax is `alpha | bravo`. `alpha` will write to the pipe when its computation is ready and then sleep in the wait queue.[69] `bravo` will then be moved to the ready queue and soon will read from its input stream.[70] The kernel will generate *software interrupts* to coordinate the piping.[70]

Signals may be classified into 7 categories.[64] The categories are:

1. when a process finishes normally.
2. when a process has an error exception.
3. when a process runs out of a system resource.
4. when a process executes an illegal instruction.
5. when a process sets an alarm event.
6. when a process is aborted from the keyboard.
7. when a process has a tracing alert for debugging.

Hardware interrupt

Input/output (I/O) devices are slower than the CPU. Therefore, it would slow down the computer if the CPU had to wait for each I/O to finish. Instead, a computer may implement interrupts for I/O completion, avoiding the need for polling or busy waiting.[71]

Some computers require an interrupt for each character or word, costing a significant amount of CPU time. Direct memory access (DMA) is an architecture feature to allow devices to bypass the CPU and access main memory directly.[72] (Separate from the architecture, a device may perform direct memory access[c] to and from main memory either directly or via a bus.)[73][d]

Input/output

Interrupt-driven I/O

This section **needs expansion**. You can help by adding to it. *(April 2022)*

When a computer user types a key on the keyboard, typically the character appears immediately on the screen. Likewise, when a user moves a mouse, the cursor immediately moves across the screen. Each keystroke and mouse movement generates an *interrupt* called *Interrupt-driven I/O*. An interrupt-driven I/O occurs when a process causes an interrupt for every character[73] or word[74] transmitted.

Direct memory access

Devices such as hard disk drives, solid-state drives, and magnetic tape drives can transfer data at a rate high enough that interrupting the CPU for every byte or word transferred, and having the CPU transfer the byte or word between the device and memory, would require too much CPU time. Data is, instead, transferred between the device and memory independently of the CPU by hardware such as a channel or a direct memory access controller; an interrupt is delivered only when all the data is transferred.[75]

If a computer program executes a system call to perform a block I/O *write* operation, then the system call might execute the following instructions:

- Set the contents of the CPU's registers (including the program counter) into the process control block.[76]
- Create an entry in the device-status table.[77] The operating system maintains this table to keep track of which processes are waiting for which devices. One field in the table is the memory address of the process control block.
- Place all the characters to be sent to the device into a memory buffer.[66]
- Set the memory address of the memory buffer to a predetermined device register.[78]
- Set the buffer size (an integer) to another predetermined register.[78]
- Execute the machine instruction to begin the writing.
- Perform a context switch to the next process in the ready queue.

While the writing takes place, the operating system will context switch to other processes as normal. When the device finishes writing, the device will *interrupt* the currently running process by *asserting* an interrupt request. The device will also place an integer onto the data bus.[79] Upon accepting the interrupt request, the operating system will:

- Push the contents of the program counter (a register) followed by the status register onto the call stack.[58]
- Push the contents of the other registers onto the call stack. (Alternatively, the contents of the registers may be placed in a system table.)[79]
- Read the integer from the data bus. The integer is an offset to the interrupt vector table. The vector table's instructions will then:

- Access the device-status table.
- Extract the process control block.
- Perform a context switch back to the writing process.

When the writing process has its time slice expired, the operating system will:[80]

- Pop from the call stack the registers other than the status register and program counter.
- Pop from the call stack the status register.
- Pop from the call stack the address of the next instruction, and set it back into the program counter.

With the program counter now reset, the interrupted process will resume its time slice.[58]

Memory management

Among other things, a multiprogramming operating system kernel must be responsible for managing all system memory which is currently in use by the programs. This ensures that a program does not interfere with memory already in use by another program. Since programs time share, each program must have independent access to memory.

Cooperative memory management, used by many early operating systems, assumes that all programs make voluntary use of the kernel's memory manager, and do not exceed their allocated memory. This system of memory management is almost never seen any more, since programs often contain bugs which can cause them to exceed their allocated memory. If a program fails, it may cause memory used by one or more other programs to be affected or overwritten. Malicious programs or viruses may purposefully alter another program's memory, or may affect the operation of the operating system itself. With cooperative memory management, it takes only one misbehaved program to crash the system.

Memory protection enables the kernel to limit a process' access to the computer's memory. Various methods of memory protection exist, including memory segmentation and paging. All methods require some level of hardware support (such as the 80286 MMU), which does not exist in all computers.

In both segmentation and paging, certain protected mode registers specify to the CPU what memory address it should allow a running program to access. Attempts to access other addresses trigger an interrupt, which causes the CPU to re-enter supervisor mode, placing the kernel in charge. This is called a segmentation violation or Seg-V for short, and since it is both difficult to assign a meaningful result to such an operation, and because it is usually a sign of a misbehaving program, the kernel generally resorts to terminating the offending program, and reports the error.

Windows versions 3.1 through ME had some level of memory protection, but programs could easily circumvent the need to use it. A general protection fault would be produced, indicating a segmentation violation had occurred; however, the system would often crash anyway.

Virtual memory

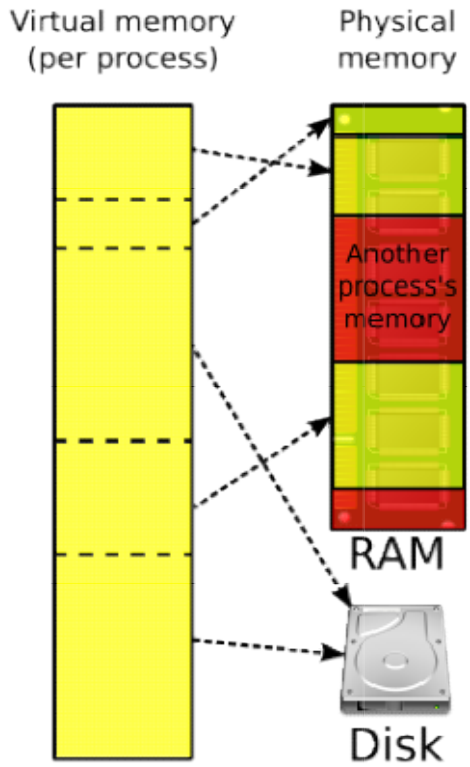

Many operating systems can "trick" programs into using memory scattered around the hard disk and RAM as if it is one continuous chunk of memory, called virtual memory.

The use of virtual memory addressing (such as paging or segmentation) means that the kernel can choose what memory each program may use at any given time, allowing the operating system to use the same memory locations for multiple tasks.

If a program tries to access memory that is not accessible[e] memory, but nonetheless has been allocated to it, the kernel is interrupted *(see § Memory management)*. This kind of interrupt is typically a page fault.

When the kernel detects a page fault it generally adjusts the virtual memory range of the program which triggered it, granting it access to the memory requested. This gives the kernel discretionary power over where a particular application's memory is stored, or even whether or not it has actually been allocated yet.

In modern operating systems, memory which is accessed less frequently can be temporarily stored on a disk or other media to make that space available for use by other programs. This is called swapping, as an area of memory can be used by

multiple programs, and what that memory area contains can be swapped or exchanged on demand.

Virtual memory provides the programmer or the user with the perception that there is a much larger amount of RAM in the computer than is really there.[81]

Concurrency

See also: Computer multitasking and Process management (computing)

Concurrency refers to the operating system's ability to carry out multiple tasks simultaneously.[82] Virtually all modern operating systems support concurrency.[83]

Threads enable splitting a process' work into multiple parts that can run simultaneously.[84] The number of threads is not limited by the number of processors available. If there are more threads than processors, the operating system kernel schedules, suspends, and resumes threads, controlling when each thread runs and how much CPU time it receives.[85] During a context switch a running thread is suspended, its state is saved into the thread control block and stack, and the state of the new thread is loaded in.[86] Historically, on many systems a thread could run until it relinquished control (cooperative multitasking). Because this model can allow a single thread to monopolize the processor, most operating systems now can interrupt a thread (preemptive multitasking).[87]

Threads have their own thread ID, program counter (PC), a register set, and a stack, but share code, heap data, and other resources with other threads of the same process Thus, there is less overhead to create a thread than a new process.[90] On single-CPU systems, concurrency is switching between processes. Many computers have multiple CPUs.[91] Parallelism with multiple threads running on different CPUs can speed up a program, depending on how much of it can be executed concurrently.[92]

File system

: Virtual file system

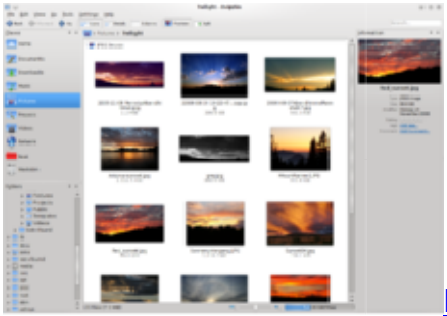

File systems allow users and programs to organize and sort files on a computer, often through the use of directories (or folders).

Permanent storage devices used in twenty-first century computers, unlike volatile dynamic random-access memory (DRAM), are still accessible after a crash or power failure. Permanent (non-volatile) storage is much cheaper per byte, but takes several orders of magnitude longer to access, read, and write.[93][94] The two

main technologies are a hard drive consisting of magnetic disks, and flash memory (a solid-state drive that stores data in electrical circuits). The latter is more expensive but faster and more durable. File systems are an abstraction used by the operating system to simplify access to permanent storage. They provide human-readable filenames and other metadata, increase performance via amortization of accesses, prevent multiple threads from accessing the same section of memory, and include checksums to identify corruption.[97] File systems are composed of files (named collections of data, of an arbitrary size) and directories (also called folders) that list human-readable filenames and other directories.[98] An absolute file path begins at the root directory and lists subdirectories divided by punctuation, while a relative path defines the location of a file from a directory.[99][100]

System calls (which are sometimes wrapped by libraries) enable applications to create, delete, open, and close files, as well as link, read, and write to them. All these operations are carried out by the operating system on behalf of the application.[101] The operating system's efforts to reduce latency include storing recently requested blocks of memory in a cache and prefetching data that the application has not asked for, but might need next.[102] Device drivers are software specific to each input/output (I/O) device that enables the operating system to work without modification over different hardware.[103][104]

Another component of file systems is a dictionary that maps a file's name and metadata to the data block where its contents are stored.[105] Most file systems use directories to convert file names to file numbers. To find the block number, the operating system uses an index (often implemented as a tree).[106] Separately, there is a free space map to track free blocks, commonly implemented as a bitmap.[106] Although any free block can be used to store a new file, many operating systems try to group together files in the same directory to maximize performance, or periodically reorganize files to reduce fragmentation.[107]

Maintaining data reliability in the face of a computer crash or hardware failure is another concern.[108] File writing protocols are designed with atomic operations so as not to leave permanent storage in a partially written, inconsistent state in the event of a crash at any point during writing.[109] Data corruption is addressed by redundant storage (for example, RAID—redundant array of inexpensive disks)[110][111] and checksums to detect when data has been corrupted. With multiple layers of checksums and backups of a file, a system can recover from multiple hardware failures. Background processes are often used to detect and recover from data corruption.[111]

Security

: Computer security

Security means protecting users from other users of the same computer, as well as from those who seeking remote access to it over a network.[112] Operating systems security rests on achieving the CIA triad: confidentiality (unauthorized users cannot access data), integrity (unauthorized users cannot modify data), and availability

(ensuring that the system remains available to authorized users, even in the event of a denial of service attack).[113] As with other computer systems, isolating security domains—in the case of operating systems, the kernel, processes, and virtual machines—is key to achieving security.[114] Other ways to increase security include simplicity to minimize the attack surface, locking access to resources by default, checking all requests for authorization, principle of least authority (granting the minimum privilege essential for performing a task), privilege separation, and reducing shared data.[115]

Some operating system designs are more secure than others. Those with no isolation between the kernel and applications are least secure, while those with a monolithic kernel like most general-purpose operating systems are still vulnerable if any part of the kernel is compromised. A more secure design features microkernels that separate the kernel's privileges into many separate security domains and reduce the consequences of a single kernel breach.[116] Unikernels are another approach that improves security by minimizing the kernel and separating out other operating systems functionality by application.[116]

Most operating systems are written in C or C++, which create potential vulnerabilities for exploitation. Despite attempts to protect against them, vulnerabilities are caused by buffer overflow attacks, which are enabled by the lack of bounds checking.[117] Hardware vulnerabilities, some of them caused by CPU optimizations, can also be used to compromise the operating system.[118] There are known instances of operating system programmers deliberately implanting vulnerabilities, such as back doors.[119]

Operating systems security is hampered by their increasing complexity and the resulting inevitability of bugs.[120] Because formal verification of operating systems may not be feasible, developers use operating system hardening to reduce vulnerabilities,[121] e.g. address space layout randomization, control-flow integrity,[122] access restrictions,[123] and other techniques.[124] There are no restrictions on who can contribute code to open source operating systems; such operating systems have transparent change histories and distributed governance structures.[125] Open source developers strive to work collaboratively to find and eliminate security vulnerabilities, using code review and type checking to expunge malicious code.[126][127] Andrew S. Tanenbaum advises releasing the source code of all operating systems, arguing that it prevents developers from placing trust in secrecy and thus relying on the unreliable practice of security by obscurity.[128]