

## Definitions:

- a program in execution
- an instance of a program running on a computer
- the entity that can be assigned to and executed on a processor
- a unit of activity characterized by the execution of a sequence of instructions, a current state, and an associated set of system instructions

## Elements of a process:

Identifier	Memory pointers
State	Context data
Priority	I/O status information
Program counter	Accounting information

## Process control block:

Contains the process elements  
Created and managed by the operating system  
Allows support for multiple processes

Computer platform consists of a collection of hardware resources

Computer applications are developed to perform some task

Inefficient for applications to be written directly for a given hardware platform

Operating system provides a convenient-to-use, feature-rich, secure, and consistent interface for applications to use

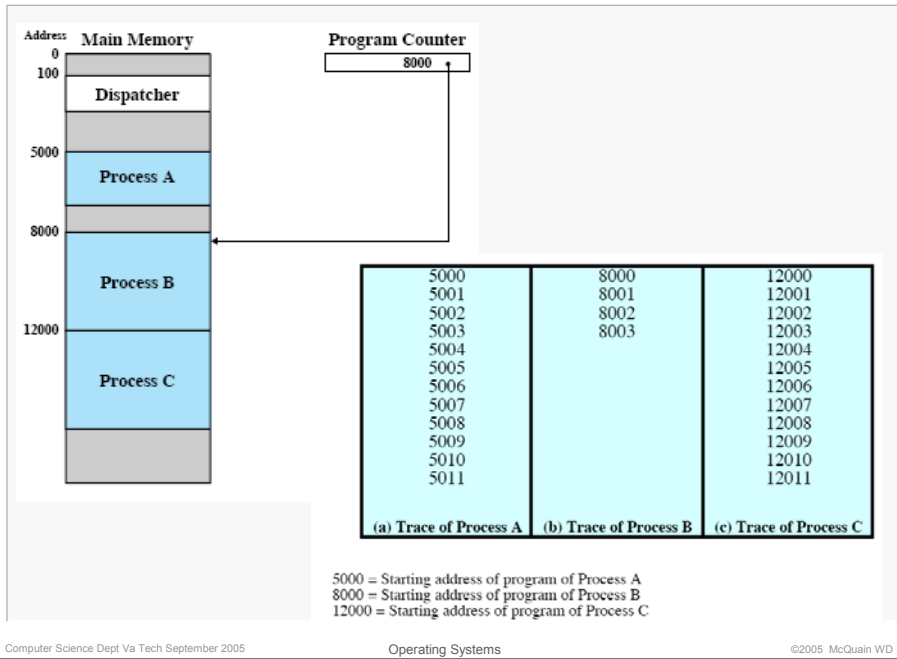
OS provides a uniform, abstract representation of resources that can be requested and accessed by application

Interleave the execution of multiple processes to maximize processor utilization while providing reasonable response time

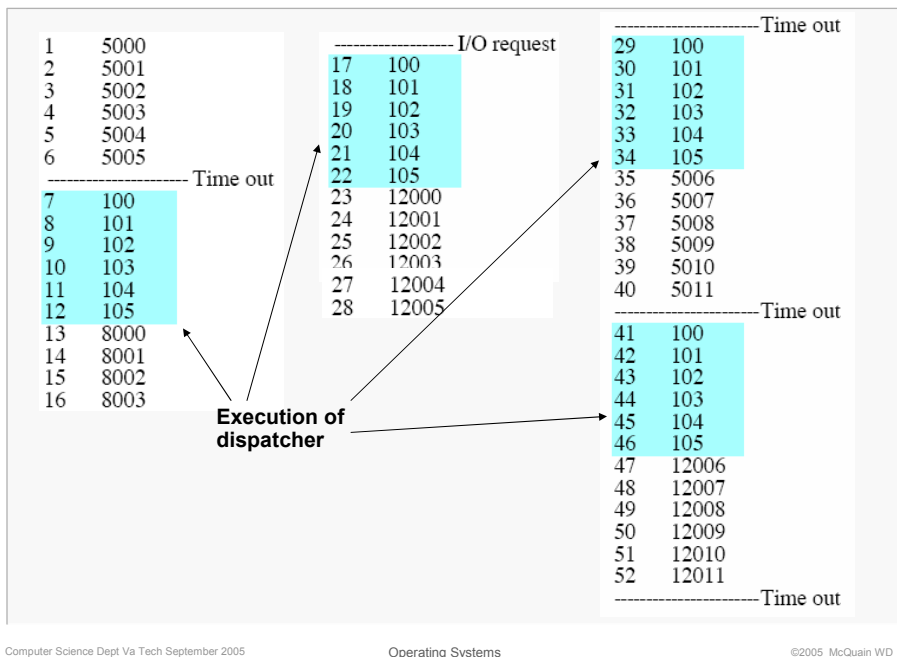
Allocate resources to processes

Support interprocess communication and user creation of processes

# Example Execution




# Process Trace




Process may be in one of two states

- Running
- Not-running



Not-running processes are kept in a scheduling queue



**Table 3.1 Reasons for Process Creation**

New batch job	The operating system is provided with a batch job control stream, usually on tape or disk. When the operating system is prepared to take on new work, it will read the next sequence of job control commands.
Interactive logon	A user at a terminal logs on to the system.
Created by OS to provide a service	The operating system can create a process to perform a function on behalf of a user program, without the user having to wait (e.g., a process to control printing).
Spawned by existing process	For purposes of modularity or to exploit parallelism, a user program can dictate the creation of a number of processes.

Table 3.2 Reasons for Process Termination

Normal completion	The process executes an OS service call to indicate that it has completed running.
Time limit exceeded	The process has run longer than the specified total time limit. There are a number of possibilities for the type of time that is measured. These include total elapsed time ("wall clock time"), amount of time spent executing, and, in the case of an interactive process, the amount of time since the user last provided any input.
Memory unavailable	The process requires more memory than the system can provide.
Bounds violation	The process tries to access a memory location that it is not allowed to access.
Protection error	The process attempts to use a resource such as a file that it is not allowed to use, or it tries to use it in an improper fashion, such as writing to a read-only file.
Arithmetic error	The process tries a prohibited computation, such as division by zero, or tries to store numbers larger than the hardware can accommodate.

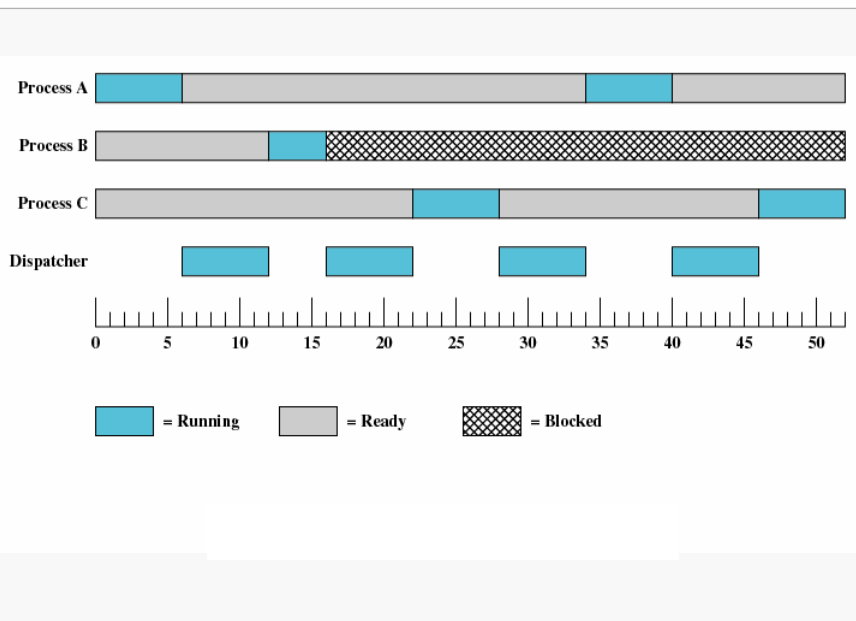
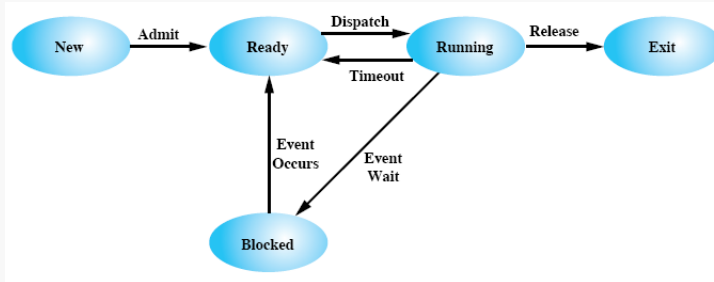
Table 3.2 Reasons for Process Termination

Time overrun	The process has waited longer than a specified maximum for a certain event to occur.
I/O failure	An error occurs during input or output, such as inability to find a file, failure to read or write after a specified maximum number of tries (when, for example, a defective area is encountered on a tape), or invalid operation (such as reading from the line printer).
Invalid instruction	The process attempts to execute a nonexistent instruction (often a result of branching into a data area and attempting to execute the data).
Privileged instruction	The process attempts to use an instruction reserved for the operating system.
Data misuse	A piece of data is of the wrong type or is not initialized.
Operator or OS intervention	For some reason, the operator or the operating system has terminated the process (for example, if a deadlock exists).
Parent termination	When a parent terminates, the operating system may automatically terminate all of the offspring of that parent.
Parent request	A parent process typically has the authority to terminate any of its offspring.

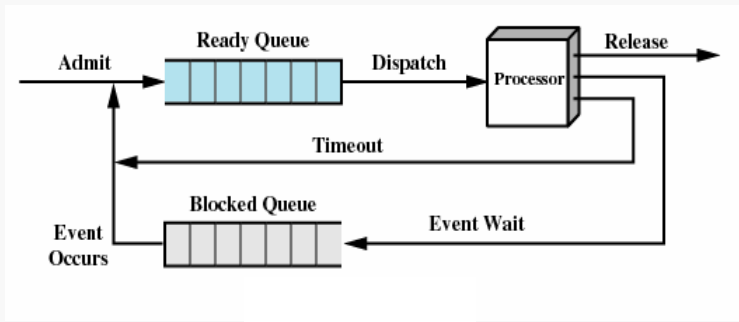
Not-running  
ready to execute

Blocked  
waiting for I/O

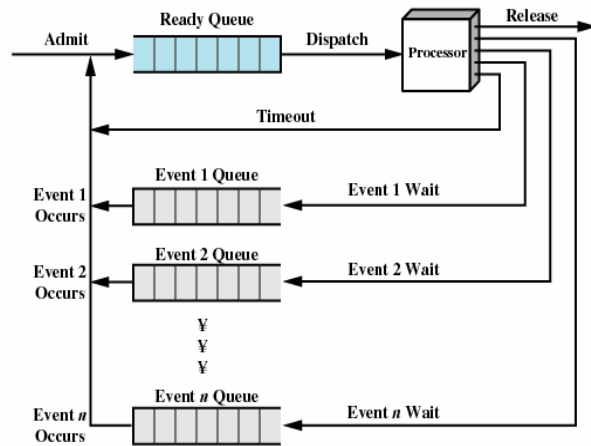
Dispatcher cannot just select the process that has been in the queue the longest because it may be blocked... leads to a more complex scheduling model:



# Using Two Queues



# Multiple Blocked Queues



## Suspended Processes


Processes 13

Processor is faster than I/O so all processes could be waiting for I/O  
 Swap these processes to disk to free up more memory  
 Blocked state becomes suspend state when swapped to disk

Two new states

Blocked/Suspend

Ready/Suspend



## Reasons for Process Suspension


Processes 14

Table 3.3 Reasons for Process Suspension

Swapping	The operating system needs to release sufficient main memory to bring in a process that is ready to execute.
Other OS reason	The operating system may suspend a background or utility process or a process that is suspected of causing a problem.
Interactive user request	A user may wish to suspend execution of a program for purposes of debugging or in connection with the use of a resource.
Timing	A process may be executed periodically (e.g., an accounting or system monitoring process) and may be suspended while waiting for the next time interval.
Parent process request	A parent process may wish to suspend execution of a descendent to examine or modify the suspended process, or to coordinate the activity of various descendents.

Information about the current status of each process and resource

Tables are constructed for each entity the operating system manages



#### Memory Tables

- allocation of main memory to processes
- allocation of secondary memory to processes
- protection attributes for access to shared memory regions
- information needed to manage virtual memory

#### I/O Tables

- I/O device is available or assigned
- status of I/O operation
- location in main memory being used as the source or destination of the I/O transfer

#### File Tables

- existence of files
- location on secondary memory
- current status
- attributes
- sometimes this information is maintained by a file management system



Where process is located

Attributes in the process control block

Program  
Data  
Stack

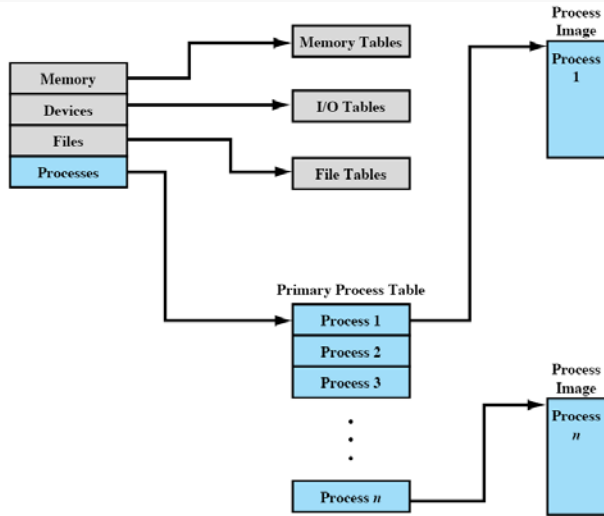


Table 3.4 Typical Elements of a Process Image

<p><b>User Data</b> The modifiable part of the user space. May include program data, a user stack area, and programs that may be modified.</p> <p><b>User Program</b> The program to be executed.</p> <p><b>System Stack</b> Each process has one or more last-in-first-out (LIFO) system stacks associated with it. A stack is used to store parameters and calling addresses for procedure and system calls.</p> <p><b>Process Control Block</b> Data needed by the operating system to control the process (see Table 3.5).</p>
--

### Process identification

#### Identifiers

Numeric identifiers that may be stored with the process control block include

- Identifier of this process
- Identifier of the process that created this process (parent process)
- User identifier

### Processor State Information

#### User-Visible Registers

A user-visible register is one that may be referenced by means of the machine language that the processor executes while in user mode. Typically, there are from 8 to 32 of these registers, although some RISC implementations have over 100.

#### Control and Status Registers

These are a variety of processor registers that are employed to control the operation of the processor. These include

- Program counter:* Contains the address of the next instruction to be fetched
- Condition codes:* Result of the most recent arithmetic or logical operation (e.g., sign, zero, carry, equal, overflow)
- Status information:* Includes interrupt enabled/disabled flags, execution mode

### Processor State Information

#### Stack Pointers

Each process has one or more last-in-first-out (LIFO) system stacks associated with it. A stack is used to store parameters and calling addresses for procedure and system calls. The stack pointer points to the top of the stack.

### Process Control Information

#### Scheduling and State Information

This is information that is needed by the operating system to perform its scheduling function. Typical items of information:

- Process state:* defines the readiness of the process to be scheduled for execution (e.g., running, ready, waiting, halted).
- Priority:* One or more fields may be used to describe the scheduling priority of the process. In some systems, several values are required (e.g., default, current, highest-allowable)
- Scheduling-related information:* This will depend on the scheduling algorithm used. Examples are the amount of time that the process has been waiting and the amount of time that the process executed the last time it was running.
- Event:* Identity of event the process is awaiting before it can be resumed

### Process Control Information

#### Data Structuring

A process may be linked to other process in a queue, ring, or some other structure. For example, all processes in a waiting state for a particular priority level may be linked in a queue. A process may exhibit a parent-child (creator-created) relationship with another process. The process control block may contain pointers to other processes to support these structures.

#### Interprocess Communication

Various flags, signals, and messages may be associated with communication between two independent processes. Some or all of this information may be maintained in the process control block.

#### Process Privileges

Processes are granted privileges in terms of the memory that may be accessed and the types of instructions that may be executed. In addition, privileges may apply to the use of system utilities and services.

### Process Control Information

#### Memory Management

This section may include pointers to segment and/or page tables that describe the virtual memory assigned to this process.

#### Resource Ownership and Utilization

Resources controlled by the process may be indicated, such as opened files. A history of utilization of the processor or other resources may also be included; this information may be needed by the scheduler.

## Processor State Information

Processes 23

### Contents of processor registers


- User-visible registers
- Control and status registers
- Stack pointers

### Program status word (PSW)

- contains status information
- Example: the EFLAGS register on Pentium machines

## Pentium II EFLAGS Register

Processes 24



- |                                 |                            |
|---------------------------------|----------------------------|
| ID = Identification flag        | DF = Direction flag        |
| VIP = Virtual interrupt pending | IF = Interrupt enable flag |
| VIF = Virtual interrupt flag    | TF = Trap flag             |
| AC = Alignment check            | SF = Sign flag             |
| VM = Virtual 8086 mode          | ZF = Zero flag             |
| RF = Resume flag                | AF = Auxiliary carry flag  |
| NT = Nested task flag           | PF = Parity flag           |
| IOPL = I/O privilege level      | CF = Carry flag            |
| OF = Overflow flag              |                            |

### User mode

Less-privileged mode

User programs typically execute in this mode

### System mode, control mode, or kernel mode

More-privileged mode

Kernel of the operating system

Assign a unique process identifier

Allocate space for the process

Initialize process control block

Set up appropriate linkages

Ex: add new process to linked list used for scheduling queue

Create or expand other data structures

Ex: maintain an accounting file

### Clock interrupt

process has executed for the maximum allowable time slice

### I/O interrupt

### Memory fault

memory address is in virtual memory so it must be brought into main memory

### Trap

error or exception occurred  
may cause process to be moved to Exit state

### Supervisor call

such as file open

Save context of processor including program counter and other registers

Update the process control block of the process that is currently in the Running state

Move process control block to appropriate queue – ready; blocked; ready/suspend

Select another process for execution

Update the process control block of the process selected

Update memory-management data structures

Restore context of the selected process


## Execution of the Operating System

Processes 29

### Non-process Kernel

Execute kernel outside of any process


Operating system code is executed as a separate entity that operates in privileged mode



### Execution Within User Processes

Operating system software within context of a user process

Process executes in privileged mode when executing operating system code




## Execution of the Operating System

Processes 30

### Process-Based Operating System

Implement operating system as a collection of system processes

Useful in multi-processor or multi-computer environment



Most of the operating system executes within the environment of a user process




Table 3.9 UNIX Process States

<b>User Running</b>	Executing in user mode.
<b>Kernel Running</b>	Executing in kernel mode.
<b>Ready to Run, in Memory</b>	Ready to run as soon as the kernel schedules it.
<b>Asleep in Memory</b>	Unable to execute until an event occurs; process is in main memory (a blocked state).
<b>Ready to Run, Swapped</b>	Process is ready to run, but the swapper must swap the process into main memory before the kernel can schedule it to execute.
<b>Sleeping, Swapped</b>	The process is awaiting an event and has been swapped to secondary storage (a blocked state).
<b>Preempted</b>	Process is returning from kernel to user mode, but the kernel preempts it and does a process switch to schedule another process.
<b>Created</b>	Process is newly created and not yet ready to run.
<b>Zombie</b>	Process no longer exists, but it leaves a record for its parent process to collect.



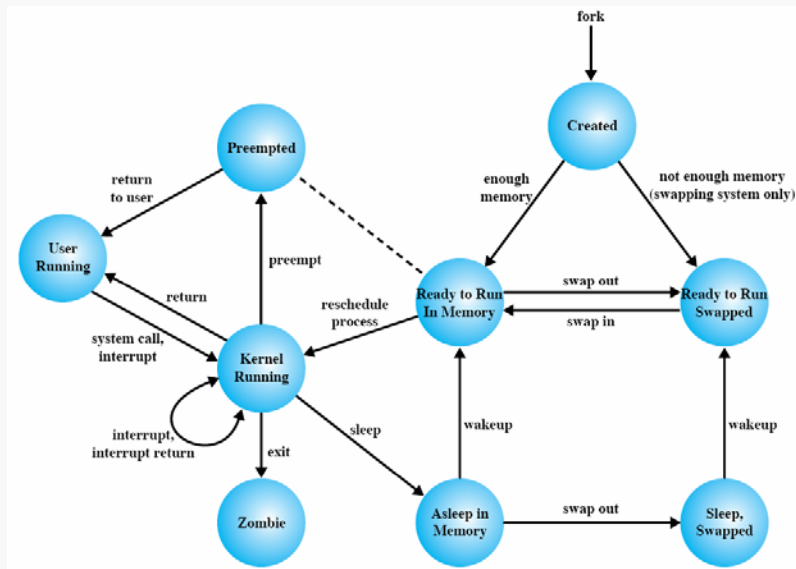


Table 3.10 UNIX Process Image

User-Level Context	
Process Text	Executable machine instructions of the program
Process Data	Data accessible by the program of this process
User Stack	Contains the arguments, local variables, and pointers for functions executing in user mode
Shared Memory	Memory shared with other processes, used for interprocess communication
Register Context	
Program Counter	Address of next instruction to be executed; may be in kernel or user memory space of this process
Processor Status Register	Contains the hardware status at the time of preemption; contents and format are hardware dependent
Stack Pointer	Points to the top of the kernel or user stack, depending on the mode of operation at the time of preemption
General-Purpose Registers	Hardware dependent
System-Level Context	
Process Table Entry	Defines state of a process; this information is always accessible to the operating system
U (user) Area	Process control information that needs to be accessed only in the context of the process
Per Process Region Table	Defines the mapping from virtual to physical addresses; also contains a permission field that indicates the type of access allowed the process: read-only, read-write, or read-execute
Kernel Stack	Contains the stack frame of kernel procedures as the process executes in kernel mode