

CS8493-OPERATING SYSTEMS

PART-A

UNIT-I

1. What are the objectives of operating systems? (MJ- 2017)

The requirements can be divided into two basic groups: User goals and System goals. Users desire that the system should be convenient and easy to use, easy to learn, reliable, safe and fast. The Operating system should be easy to design, implement, and maintain. Also it should be flexible, reliable, error free and efficient. These are some of the requirements, which are vague and have no general solution.

2. What are the advantages of peer to peer systems over client-server systems?(MJ-2016)

Easy and simple to set up only requiring a hub or a switch to connect all computers together. You can access any file on the computer as-long as it is set to a shared folder. If one computer fails to work all the other computers connected to it still continue to work.

3. What is the purpose of system programs? (MJ-2016)

System programs can be thought of as bundles of useful **system** calls. They provide basic functionality to users so that users do not need to write their own **programs** to solve common problems.

4. Compare and contrast DMA and Cache memory. (N/D 2015)

DMA is a hardware device that can move to/from memory without using CPU instructions. Cache stores memory to somewhere closer to where you can access so you don't have to spend as much time for the I/O. DMA provides this capability to carry out memory specific operations with minimal processor intervention. When any I/O device needs a memory access. It sends a DMA request (in form of interrupt) to the processor. The processor initiates the transfer by providing appropriate grant signals to the data bus. And passes the control to the DMA controller which controls the rest of the data transfer and transfers the data directly to I/O device.

5. Write the difference between Batch Systems and Time Sharing systems. (N/D 2015)

Batch systems are quite appropriate for executing large jobs that need little interaction. The user can submit jobs and return later for the results. It is not necessary to wait while the job is processed. Operators batched together jobs with similar needs and ran them through the computer as a group. Time-sharing or multitasking is a logical extension of multiprogramming. It allows many users to share the computer simultaneously.

6. Do timesharing differ from Multiprogramming? If so, How? (MJ- 2015)

Time-sharing or multitasking is a logical extension of multiprogramming. It allows many users to share the computer simultaneously. The CPU executes multiple jobs by switching among them, but the switches occur so frequently that the users can interact with each program while it is running.

7. Why API's need to be used rather than system calls? (MJ- 2015)

An **API** is a generic term for a way to interface with an existing library / service etc... Whereas a **system call** is about an application calling down the stack to the underlying resources and such, usually via a standard API in its own right. System calls provide the interface between a process and the Operating system.

8. Can multiple user level threads achieve better performance on a multiprocessor system than a single processor system? Justify your answer.(MJ-2014)

We assume that user-level threads are not known to the kernel. In that case, the answer is because the scheduling is done at the process level. On the other hand, some OS allows user-level threads to be assigned to different kernel level processes for the purposes of scheduling. In this case the multithreaded solution could be faster.

9. Mention the circumferences that would a user be better off using time-sharing system rather than a PC or a single user work station? (MJ-2014)

A user is better off under three situations: when it is cheaper, faster, or easier. For example:

- When the user is paying for management costs and the costs are cheaper for a time-sharing system than for a single-user computer.
- When running a simulation or calculating that takes too long to run on a single PC or workstation.
- When a user is travelling and doesn't have laptop to carry around, they can connect remotely to a time-shared system and do their work.

10. What is meant by Mainframe Systems? [MJ-15]

Mainframe systems are the first computers developed to tackle many commercial and scientific applications. These systems are developed from the batch systems and then multiprogramming system and finally time sharing systems

11. What is meant by Multiprogramming? [ND-11]

Several users simultaneously compete for system resources (i.e) the job currently waiting for I/O will yield the CPU to another job which is ready to do calculations, if another job is waiting. Thus it increases CPU utilization and system throughput.

12. What are the advantages of Multiprogramming? [ND-11],[MJ-13]

- Increased System Throughput
- Increased CPU utilization

13. What are Operating Services? [MJ-12][MJ-15]

Normally, an operating system provides certain services to programs and to the users of those programs. Some of them are: Program Execution. I/O operations File-system manipulation Communications Error Detection

14. What is System Programs?[MJ-11]

System programs provide a convenient environment for program development and execution. Some of these programs are user interfaces to system calls and others are more complex. Some of them are: File Management Status Information File modification Programming Language support Program loading, execution and communication.

15. What is meant by System Calls? [MJ-11],[MJ-14]

The System Calls acts as a interface to a running program and the Operating system. These system calls available in assembly language instructions.

16. What is meant by Batch Systems?

In this, operators batched together jobs with similar needs and ran through the computer as a group .The operators would sort programs into batches with similar requirements and as system become available, it would run each batch.

17. What are Multiprocessor System?

Multiprocessor systems have systems more than one processor for communication, sharing the computer bus, the memory, clock & peripheral devices

18. What are the Components of a Computer System?

- Application Programs System
- Program Operating System
- Computer Hardware

19. What are the advantages of multiprocessors??

- Increased throughput
- Economy of scale
- Increased reliability

20. What is meant by Real time system?

Real time systems are systems that have their iND-built characteristics as supplying immediate response. In real time system, each process is assigned a certain level of priority according to the relative importance of the events to be processed

UNIT-II

1. “Priority inversion is a condition that occurs in real time systems where a low priority process is starved because higher priority processes have gained hold of the CPU”- Comment on this statement. (MJ- 2017)

Priority inversion is a problem that occurs in concurrent processes when low-priority threads hold shared resources required by some high-priority threads, causing the high priority-threads to block indefinitely. This problem is enlarged when the concurrent processes are in a real time system where high- priority threads must be served on time.

2. Differentiate single threaded and multi threaded processes. (MJ- 2017)

A thread is a simple flow of instruction. An application can be **single threaded** (so imagine it as a single line going from the entry point of the application to its end) or **multi-threaded** (imagine a tree: the whole app starts from 1 point, than it branches out more and more).

3. Under what circumstances is user level threads is better than the kernel level threads? (MJ-2016) What are the differences between user level threads and kernel level threads? Under circumstances is one type better than the other? (N/D 2015)

User-level threads are threads that the OS is not aware of. They exist entirely within a process, and are scheduled to run within that process's time slices. The OS is aware of kernel-level threads. Kernel threads are scheduled by the OS's scheduling algorithm, and require a "lightweight" context switch to switch between (that is, registers, PC, and SP must be changed, but the memory context remains the same among kernel threads in the same process). User-level threads are much faster to switch between, as there is no context switch; further, a problem domain-dependent algorithm can be used to schedule among them. CPU-bound tasks with interdependent computations, or a task that will switch among threads often, might best be handled by user-level threads. Kernel-level threads are scheduled by the OS, and each thread can be granted its own time slices by the scheduling algorithm. The kernel scheduler can thus make intelligent decisions among threads, and avoid scheduling processes which consist of entirely idle threads (or I/O bound threads). A task that has multiple threads that are I/O bound, or that has many threads (and thus will benefit from the additional time slices that kernel threads will receive) might best be handled by kernel threads. Kernel-level threads require a system call for the switch to occur; user-level threads do not.

4. What is the meaning of the term busy waiting? (MJ-2016)

When a process is in its critical section, any other process that tries to enter its critical section must loop continuously in the entry code. This is called as busy waiting and this type of semaphore is also called a spinlock, because the process while waiting for the lock.

5. What is the concept behind strong semaphore and spinlock? (N/D 2015)

- A semaphore 'S' is a synchronization tool which is an integer value that, apart from initialization, is accessed only through two standard atomic operations; wait and signal.
- Semaphores can be used to deal with the ND-process critical section problem. It can be also used to solve various synchronization problems.

The classic definition of 'wait'

```
Wait (S)
{
while (S<=0)
S--;
}
```

The classic definition of 'signal'

```
Signal (S)
{
S++;
}
```

When a process is in its critical section, any other process that tries to enter its critical section must loop continuously in the entry code. This is called as busy waiting and this type of semaphore is also called a spinlock, because the process while waiting for the lock.

6. List out the data fields associate with process control block? (MJ- 2015)

Each process is represented in the operating system by a process control block also called a task control block. It contains many pieces of information associated with a specific process. It simply acts as a repository for any information that may vary from process to process. It contains the following information:

- Process state
- Program counter
- CPU registers
- CPU-scheduling information
- Memory-management information
- Accounting Information
- I/O status information

7. Define the term “Dispatch latency”. (MJ- 2015)

The time taken by the dispatcher to stop one process and start another running is known as dispatch latency.

8. What is the resident set and working set of a process?(N/D 2014)

Resident set is that portion of the **process** image that is actually in real-memory at a particular instant.

Working set is that subset of **resident set** that is actually needed for execution.

9. What resources are used when a thread created? (N/D 2014)

Because a thread is smaller than a process, thread creation typically uses fewer resources than process creation. Creating a process requires allocating a process control block (PCB), a rather large data structure. The PCB includes a memory map, list of open files, and environment variables. Allocating and managing the memory map is typically the most time-consuming activity. Creating either a user or kernel thread involves allocating a small data structure to hold a register set, stack, and priority.

10. Differentiate between pre-emptive and noND-pre-emptive scheduling. (N/D 2014)

Under non - preemptive scheduling once the CPU has been allocated to a process, the process keeps the CPU until it releases the CPU either by terminating or switching to the waiting state.

Preemptive scheduling can preempt a process which is utilizing the CPU in between its execution and give the CPU to another process.

11. What can the operating system do to recover from deadlock? (N/D 2014)

Deadlock detection and **recovery** - Abort a process or preempt some resources when **deadlocks** are detected. It can kill processes until the deadlock is resolved. Or it could preempt the resources and simply take them away from the process. It could also rollback to an earlier state before the deadlock and hope that it doesn't happen again.

12. Write the four situations under which CPU scheduling decisions takes place. (MJ-2014)

- CPU scheduling decisions take place under one of four conditions:
- When a process switches from the running state to the waiting state, such as for an I/O request or invocation of the wait () system call.
- When a process switches from the running state to the ready state, for example in response to an interrupt.
- When a process switches from the waiting state to the ready state, say at completion of I/O or a return from wait ().
- When a process terminates.

13. Show that mutual exclusion may be violated if the signal and wait operations are not executed automatically. (MJ-2014)

A wait operation atomically decrements the value associated with a semaphore. If two wait operations are executed on a semaphore when its value is 1, if the two operations are not performed atomically, then it is possible that both operations might proceed to decrement the semaphore value, thereby violating mutual exclusion.

14. What is FPCB? [ND-12]

FPCB is a data structure containing certain important information about the process including the following:
Current state of the process
Unique identification of the process
A pointer to the process's parent
A pointer to the process's child.
The process's priority
Pointers to locate the process's memory and to allocate resources.

15. Define Mutual Exclusion. [ND-11]

Each process accessing the shared data excludes all others from doing simultaneously called as Mutual Exclusion.

16. What are the conditions that must hold for Deadlock Prevention? [MJ-11][MJ-12][MJ-13] [ND-13][MJ-15]

- Mutual Exclusion Condition
- Hold and Wait Condition
- No Pre-emption condition
- Circular Wait Condition.

17. What are the options for breaking a Deadlock? [MJ-11], [MJ-12],[MJ-15]

Simply abort one or more process to break the circular wait. Preempt some resources from one or more of the deadlocked processes.

- Resource-Allocation Graph Algorithm
- Banker's Algorithm
 - Safety Algorithm
 - Resource-Request Algorithm

18. What is meant by Counting Semaphore? [ND-11], [ND-12]

A Counting Semaphore is a semaphore whose integer value that can range between 0 & 1.

19. What is meant by Binary Semaphore? [ND-11],[ND-12]

A Binary Semaphore is a semaphore with an integer value that can range between 0 and 1. It can be simpler to implement than a counting semaphore, depending on the underlying hardware architecture.

20. What is meant by Indefinite Blocking or Starvation? [MJ-15]

Indefinite Blocking is a situation where process waits indefinitely within the semaphore. This may occur if we add and remove processes from the list associated with a semaphore in LIFO order.

21. What is meant by CPU Scheduler?[MJ-11]

When the CPU becomes idle, the operating system must select one of the processes in the ready queue to be executed. This selection process is carried out by the CPU Scheduler.

22. What is meant by CPU Scheduling? [MJ-11]

The process of selecting among the processes in memory that are ready to execute and allocates the CPU to one of them is called as CPU Scheduling.

23. What are the types of Scheduling available? [MJ-12]

- Preemptive Scheduling
- Non - preemptive Scheduling
- Priority Scheduling

24. What is Preemptive Scheduling? [ND-14][MJ-15]

A Scheduling discipline is Pre-emptive if the CPU can be taken away before the process completes.

25. What is non - Preemptive Scheduling? [ND-14][MJ-15]

A Scheduling discipline is non pre-emptive if once a process has been given the CPU, the CPU cannot be taken away from the process.

26. What is known as Resource Reservation in Real time Scheduling? [ND-12]

The Scheduler either admits the process, guaranteeing that the process will complete ontime or rejects the request as impossible. This is known as Resource Reservation.

27. What is meant by First Come, First Served Scheduling? [MJ-14]

In this Scheduling, the process that requests the CPU first is allocated the CPU first. This Scheduling algorithm is Non Pre-emptive.

UNIT-III

1. What is the difference between a user-level instruction and a privileged instruction? Which of the following instructions should be privileged and only allowed to execute in kernel mode? (MJ- 2017)

- a) **Load a value from a memory address to a general-purpose register.**
 - b) **Set a new value in the program counter (PC) register.**
 - c) **Turn off interrupts.**
- Machines have two kinds of instructions
 - “Normal” instructions, e.g., add, sub, etc.
 - “privileged” instructions, e.g., initiate I/O switch state vectors or contexts load/save from protected memory
 - **Load a value from a memory address to a general-purpose register- privileged and only allowed to execute in kernel mode**

2. Will optimal page replacement algorithm suffer from Belady's anomaly? Justify your answers. (MJ- 2017)

Yes, a page replacement algorithm suffers from Belady's anomaly when it is not a stack algorithm. A Stack algorithm is one that satisfies the inclusion property.

3. Name two differences between logical and physical addresses. (MJ-2016)

The set of all logical addresses generated by a program is called a logical address space; the set of all physical addresses corresponding to these logical addresses is a physical address space.

4. How does the system detect thrashing? (MJ-2016)

Thrashing is a situation when the degree of multitasking is not supported by the number of frames supplied by the target machine. This results in a large number of page faults and therefore a high number of swapping to and from disk. The systems then spend more time swapping than it does processing. A high number of page faults are generally a good indicator of thrashing. To recover a number of processed needs to be suspended until other complete, reducing the degree of multitasking.

5. Define demand paging in memory management. What are the steps required to handle a page fault in demand paging? (N/D 2015)

Demand paging (as opposed to anticipatory paging) is a method of virtual memory management. In a system that uses demand paging, the operating system copies a disk page into physical memory only if an attempt is made to access it and that page is not already in memory (i.e., if a page fault occurs). It follows that a process begins execution with none of its pages in physical memory, and many page faults will occur until most of a process's working set of pages is located in physical memory. This is an example of a lazy loading technique.

6. What do you mean by thrashing? (MJ- 2015)

Thrashing is caused by under allocation of the minimum number of pages required by a process, forcing it to continuously page fault.

7. Mention the significance of LDT and GDT in segmentation. (MJ- 2015)

The *GDT* can hold things other than segment descriptors as well. Every 8-byte entry in the *GDT* is a descriptor, but these can be Task State Segment (or TSS) descriptors, Local Descriptor Table (LDT) descriptors, or Call Gate descriptors. The last ones, Call Gates, are particularly important for transferring control between x86 privilege levels although this mechanism is not used on most modern operating systems.

There is also an LDT or *Local Descriptor Table*. The LDT is supposed to contain memory segments which are private to a specific program, while the GDT is supposed to contain global segments.

8. What is the maximum file size supported by a file system with 16 direct blocks, single, double, and triple indirection? -The block size is 512 bytes. Disk block. Numbers can be stored in 4 bytes. (N/D 2014)

$$512 * (16 + 128 + 128 * 128 + 128 * 128 * 128) = 2,113,$$

$$580 * 512 = 1,082,204,$$

$$160 \text{ bytes} = 1,056,$$

$$840 \text{ kbytes} = 1,032 \text{ MB} \sim 1 \text{ gig}$$

9. List the steps needed to perform page replacement. (N/D 2014)

1. Determine which page is to be removed from the memory
2. Perform a page-out operation
3. Perform page-in operation

10. Consider a logical address space of eight pages of 1024 words each, mapped on to a physical memory of 32 frames. Find the number of bits in the logical address and the physical address. (MJ-2014)

Addressing within a 1024-word page requires 10 bits because $1024 = 2^{10}$. Since the logical address space consists of $8 = 2^3$ pages, the logical addresses must be $10 + 3 = 13$ bits. Similarly, since there are $32 = 2^5$ physical pages, physical addresses are $5 + 10 = 15$ bits long.

Physical Address (P = page number bits)

P P P P P - - - - -

Logical Address (P = page number bits)

P P P - - - - -

11. Define virtual memory. (MJ-2014)

Virtual memory is a **memory** management capability of an OS that uses hardware and software to allow a computer to compensate for physical **memory** shortages by temporarily transferring data from random access **memory** (RAM) to disk storage.

12. What is meant by Swapping? [ND-13]

Process of bringing in each process in its entirety, running it for a while and then putting it back on the disk.

13. What is meant by External Fragmentation and Internal Fragmentation? [ND-11] [ND-13]

External Fragmentation exists when enough total memory space exists to satisfy a request, but it is not contiguous and storage is fragmented into a large number of small holes. The memory allocated to a process may be slightly larger than the requested memory. The difference between these two numbers is called as Internal Fragmentation.

14. What is meant by Paging? Give its advantages. [ND-11]

Paging is a Memory-management scheme that permits the physical -address space of a process to be NoND-contiguous.

Advantages:

Avoids the considerable problem of fitting the varying -sized memory chunks onto the baking store
Fragmentation problems are also prevalent baking store, except that access is much slower, so compaction is impossible.

15. What is TLB and Hit-Ratio? [MJ-11]

Translation Look aside Buffer (TLB) is a small, special and fast cache which is associated with high speed memory. The Percentage of times that a particular page number is found in the Translation Look aside Buffer (TLB) is called as Hit- Ratio.

16. What is meant by Demand Paging? [MJ-12]

Whenever the CPU tries to fetch the first instruction, it gets a page fault causing the OS to bring in the page containing that instruction. Thus the pages are loaded only on demand is called as Demand Paging.

17. What is meant by Thrashing? [MJ-11][MJ-15]

A Program which is causing page faults every few instructions to occur is called as Thrashing.

18. What is meant by Segmentation?

Segmentation is a memory-management scheme that supports the user-view memory. Blocks of different size is called as Segments and its associative virtual storage Organization is called as Segmentation.

19. What do you mean by Best Fit?

Best fit allocates the smallest hole that is big enough. The entire list has to be searched, unless it is sorted by size. This strategy produces the smallest leftover hole.

20. What do you mean by First Fit?

First fit allocates the first hole that is big enough. Searching can either start at the beginning of the set of holes or where the previous first-fit search ended. Searching can be stopped as soon as a free hole that is big enough is found.

UNIT-IV

1. Differentiate file and directory. (MJ- 2017)

All the data on your hard drive consists of files and folders. The basic difference between the two is that files store data, while folders store files and other folders. The folders, often referred to as directories, are used to organize files on your computer. The folders themselves take up virtually no space on the hard drive. Files, on the other hand, can range from a few bytes to several gigabytes. They can be documents, programs, libraries, and other compilations of data.

2. Why rotational latency is usually not considered in disk scheduling? (MJ-2016)

Most disks do not export their rotational position information to the host. Even if they did, the time for this information to reach the scheduler would be subject to imprecision and the time consumed by the scheduler is variable, so the rotational position information would become incorrect. Further, the disk requests are usually given in terms Mass-Storage Structure of logical block numbers, and the mapping between logical blocks and physical locations is very complex.

3. How does DMA increase system concurrency? (MJ-2016)

DMA increases system concurrency by allowing the CPU to perform tasks while the **DMA system** transfers data via the **system** and memory buses. Hardware design is complicated because the **DMA** controller must be integrated into the **system** and the **system** must allow the **DMA** controller to be a bus master.

- 4. List out the major attributes and operations of a file. (MJ- 2015) List the attributes of a file. (MJ-2014)**
- **Read** - Designated as an "r"; allows a file to be read, but nothing can be written to or changed in the file.
 - **Write** - Designated as a "w"; allows a file to be written to and changed.
 - **Execute** - Designated as an "x"; allows a file to be executed by users or the operating system.

5. What is HSM? Where it is used? (MJ- 2015)

A hardware security module (**HSM**) is a physical computing device that safeguards and manages digital keys for strong authentication and provides crypto processing. These modules traditionally come in the form of a plug-in card or an external device that attaches directly to a computer or network server.

6. What file access pattern is -particularly suited to chained file allocation on disk? (N/D 2014)

Linked List of Free Blocks

A naive implementation would simply link the free blocks together and just keep a pointer to the head of the list. This simple scheme has poor performance since it requires an extra I/O for every acquisition or return of a free block.

7. What file allocation strategy is most appropriate for random access files? (N/D 2014)

Contiguous allocation as you do not have to follow a linked list to data, it can be randomly accessed.

8. Compare bitmap-based allocation of blocks on disk with a free block list. (N/D 2014)

A free block-list after a period of time during which numerous allocations and deal locations take place will result in a random list of free space. This does not support contiguous allocation as the free list will force the FS to place data randomly. Bit-map-based allocation maintains the order of the free blocks, easily supporting contiguous block allocation, but sacrifices time required to find free space and to release space, they are no longer order 1 as the bit-map must be traversed. The benefits far out way those small overheads when accessing files stored.

9. What is an I/O buffer? (N/D 2014)

The process of storing data in memory area called buffers while data is being transferred between two devices or between a device and an application.

10. What are the operations performed in a Directory? [ND-12]

- 1) Create
- 2) Delete
- 3) Opendir
- 4) Closedir
- 5) Readdir
- 6) Rename
- 7) Link
- 8) Unlink

11. What is meant by Boot Control block? [MJ-11]

The Block which contains information needed by the system to boot an operating system from that partition is called as Boot Control Block.

12. What is meant by Partition Control Block?[MJ-11]

The Block which contains partition details such as the number of blocks in that partition, size of the blocks, free -block count and free - block pointers is called as partition control Block.

13. What is meant by Free Space List? [ND-12]

The list which maintains/records all free disk block which means blocks that are not allocated to some file or Directory.

14. What is meant by Free Space List? [ND-12]

The list which maintains/records all free disk block which means blocks that are not allocated to some file or Directory.

15. What is Double Buffering? [ND-13]

It is a process in which the first buffer receives and holds the records generated by the running process until it becomes full. Thus the process continues to deposit the generated records in first buffer.

16. Mention few Page Replacement Strategies. [ND-14]

Optimal Page Replacement FIFO Page Replacement LRU Page replacement MFU Page Replacement LFU Page Replacement Random Page Replacement

17. What is meant by Global Replacement and Local Replacement? [ND-12]

Global Page Replacement allows a process to select a replacement frame from the set of all frames, even if that frame is currently allocated to some other process. Local Replacement requires that each process select from only its own set of allocated frames. Here the number of frames allocated to a process doesn't change.

18. What is meant by Working Set? [MJ-14] [ND-14]

A Working Set is defined as the collection of pages a process is actively referencing.

19. What are File Attributes? [MJ-14]

- Identifier Type,
- Size Location,
- Protection Time,
- Date & User Identification

20. What is meant by Seek Time? [MJ-12]

It is the time taken for the disk arm to move the heads to the cylinder containing the desired sector.

21. What is meant by Rotational Latency? [MJ-12][MJ-13]

It is defined as the additional time waiting for the disk to rotate the desired sector to the disk head.

22. What is meant by Swap-Space Management? [ND-12]

It is a low-level task of the operating system. Efficient management of the swap space is called as Swap space management. This Swap space is the space needed for the entire process image including code and Data segments.

23. What is meant by Disk Scheduling? [ND-12] [MJ-14]

Disk scheduling is a process of allocation of the disk to one process at a time. In multiprogrammed system, many processes try to read or write the records on disks at the same time. To avoid disk arbitration, it is necessary.

24. Why Disk Scheduling necessary?[ND-12]

To avoid Disk arbitration which occurs when many processes try to read or write the records on disks at the same time, Disk Scheduling is necessary.

25. What are the characteristics of Disk Scheduling? [ND-12]

- 1) Throughput
- 2) Mean Response Time
- 3) Variance of Response time

26. What are the different types of Disk Scheduling? [ND-12] [MJ-14]

- (i) SSTF Scheduling
- (ii) FCFS Scheduling
- (iii) SCAN Scheduling
- (iv) C-SCAN Scheduling
- (v) LOOK Scheduling.

UNIT-V

1. Mention any two features of Linux file systems. (MJ- 2017)

- Portable
- Open Source
- Multi-User
- Multiprogramming
- Hierarchical File System
- Shell
- Security

2. Enlist the advantages of using kernel modules in Linux. (MJ- 2017)

- Loading and unloading a module is much more flexible and faster than recompiling a kernel and rebooting.
- You can try different options each time you load a module. Most drivers that handle hardware will take options for I/O addresses, IRQ or DMA numbers, plus more esoteric options like full or half duplex. When you have problems getting a card to run correctly, the ability to try different options can save hours.
- Makes it easier to maintain multiple machines on a single kernel base.

3. State the components of the Linux system. (MJ-2016)

- (i).Kernel
- (ii).System libraries
- (iii).System utilities

4. Why virtualization is required? (N/D 2015)

Virtualization essentially means to create multiple, logical instances of software or hardware on a single physical hardware resource. This technique simulates the available hardware and gives every application running on top of it, the feel that it is the unique holder of the resource. The details of the virtual, simulated environment are kept transparent from the application. Organizations may use this technique as illustrated by the video to actually do away with many of their physical servers and map their function onto one robust, evergreen physical server. The advantage here is the reduced cost of maintenance and reduced energy wastage which is not very surprising.

5. What is the responsibility of kernel in LINUX operating system? (MJ- 2015)

Kernel is responsible for maintaining all the important abstractions of the operating system including such things as virtual memory and processes.

6. What is the significance of microkernel design?(AOS-MJ-2015)

- A microkernel is a piece of software or even code that contains the near-minimum amount of functions and features required to implement an operating system.
- It provides the minimal number of mechanisms, just enough to run the most basic functions of a system, in order to maximize the implementation flexibility so it allows for other parts of the OS to be implemented efficiently since it does not impose a lot of policies.

7. What are the components of Kernel modules? (AOS-MJ-2015)

The Linux kernel is a monolithic kernel i.e. it is one single large program where all the functional components of the kernel have access to all of its internal data structures and routines. The alternative to this is the micro kernel structure where the functional pieces of the kernel are broken out into units with strict communication mechanism between them.

8. What do you mean by Virtual File Systems? (AOS-MJ-2015) Identify the two important functions of Virtual File System (VFS) layer in the concept of file system implementation. (N/D 2015)

A virtual file system (VFS) is programming that forms an interface between an operating system's kernel and a more concrete file system.

The VFS serves as an abstraction layer that gives applications access to different types of file systems and local and network storage devices. For that reason, a VFS may also be known as a *virtual file system switch*. It also manages the data storage and retrieval between the operating system and the storage sub-system.

9. What is meant by Data Striping?.

Data Striping means splitting the bits of each byte across multiple disks .It is also called as Bit -level Striping.

13. What are the main supports for the Linux modules?

- (i). Module Management
- (ii).Driver Registration.
- (iii).Conflict Resolution mechanism.

16. What is meant by Personality?

Process Personalities are primarily used by emulation libraries to request that system call be compatible with certain versions of UNIX.

17. What is meant by Buffer cache?

It is the kernel's main cache for block-oriented devices such as disk drives and is the main mechanism through which I/O to these devices is performed.

19. What is meant by Kernel in Linux system?

Kernel is responsible for maintaining all the important abstractions of the operating system including such things as virtual memory and processes.

20. What is meant by System Libraries?

System Libraries define a standard set of functions through which applications can interact with the kernel and that implement much of the operating -system functionality that doesn't need the full privileges of kernel code.