READY TO LEARN
RLA
LEARN TO ACHIEVE

TAMILNADU PUBLIC SERVICE COMMISSION

TNPSC-AE

ASSISTANT ENGINEER

YouTube @ RLA ACADEMY

2000 MULTIPLE CHOICE QUESTIONS

COMPUTER SCIENCE, ELECTRONICS & ELECTRICAL ENGINEERING, INFORMATION TECHNOLOGY

ENGLISH MEDIUM

CODE : 554

BASED ON SYLLABUS

10 UNITS COVERED

USEFUL FOR

AE | SSC JE | RRB JE

UNIT V: OPERATING SYSTEMS AND CLOUD TECHNOLOGIES

QUESTION

1. In a demand–paged virtual memory system, the Effective Access Time (EAT) is represented by a formula involving memory access time (m), page fault rate (p), and page fault service time (f). Which expression correctly represents the Effective Access Time?

A) EAT = $(1 - p)m + pf$

B) EAT = $m + pf$

C) EAT = $p(m + f)$

D) EAT = $(m \times f)/p$

2. A DMA controller transfers a block of 4096 bytes directly from an SSD to main memory over a 32–bit system bus. Assuming one bus transfer moves 32 bits of data and no wait states occur, how many bus transfer cycles are required to complete the DMA operation?

A) 512

B) 1024

C) 2048

D) 4096

3. A cloud data center is deploying hundreds of virtual machines on a multicore server. The administrator requires a software layer that can efficiently allocate CPU, memory, storage, and network resources among multiple guest operating systems while maintaining isolation. Which component is the most appropriate choice?

A) Device Driver

B) Hypervisor

C) Bootstrap Loader

D) Shell Interpreter

4. Which operating system component acts as an interface between user applications and hardware resources through controlled service requests?

A) Cache Controller
B) System Call Interface
C) DMA Controller
D) Bootstrap Program

5. Which of the following is NOT a characteristic objective of an operating system?

A) Resource allocation
B) Convenience for users
C) Program execution management
D) Increasing transistor density in processors

6. A system has a page size of 4 KB and logical address space of 64 MB. How many pages are present in the logical address space?

A) 8192
B) 16384
C) 32768
D) 65536

7. Which property of semaphores allows synchronization among concurrent processes?

A) Mutual exclusion capability
B) Data compression capability
C) Address translation capability
D) Instruction decoding capability

8. Identify the correct sequence during interrupt processing.

A) Save Context → Execute ISR → Restore Context → Resume Process
B) Execute ISR → Save Context → Resume Process → Restore Context
C) Restore Context → Execute ISR → Save Context → Resume Process
D) Resume Process → Save Context → Execute ISR → Restore Context

9. A real–time industrial controller requires deterministic scheduling and bounded response times. Which scheduling algorithm is most suitable?

A) First Come First Served
B) Round Robin
C) Earliest Deadline First
D) Random Scheduling

10. Which term refers to a process currently being executed by the CPU?

A) Ready Process
B) Blocked Process
C) Running Process
D) Zombie Process

11. Assertion (A): Multithreading improves responsiveness in applications.
Reason (R): Multiple threads can execute concurrently while sharing process resources.

A) Both A and R are true and R explains A
B) Both A and R are true but R does not explain A
C) A is true but R is false
D) A is false but R is true

12. A CPU utilization graph shows utilization increasing from 30% to 90% after introducing multiprogramming. What is the most likely reason?

A) Increased context switching only
B) Better overlap of CPU and I/O activities
C) Reduced memory size
D) Increased interrupt latency

13. A Round Robin scheduler uses a time quantum of 5 ms. A process requires 23 ms CPU time. Assuming no overhead, how many CPU allocations will the process receive?

A) 3
B) 4
C) 5
D) 6

14. Which formula represents CPU utilization?

A) Busy Time / Total Time × 100
B) Idle Time / Busy Time × 100
C) Turnaround Time / Waiting Time × 100
D) Throughput / Response Time

15. A cache memory has an access time of 5 ns and main memory has 80 ns access time. If cache hit ratio is 90%, find average memory access time.

A) 12.5 ns
B) 15 ns
C) 20 ns
D) 25 ns

16. Which storage technology is most suitable for implementing a high-speed paging device in cloud servers?

A) Magnetic Tape
B) Optical Disk
C) NVMe SSD
D) Punch Card

17. What is the primary role of a monitor in process synchronization?

A) Memory allocation
B) Encapsulating shared data and synchronization operations
C) Disk scheduling
D) Address translation

18. Which of the following is NOT a deadlock necessary condition?

A) Mutual Exclusion
B) Hold and Wait
C) Circular Wait
D) Priority Inversion

19. A 32-bit system uses byte addressing. What is the maximum directly addressable memory space?

A) 2 GB
B) 4 GB
C) 8 GB
D) 16 GB

20.Which property distinguishes multicore processors from multiprocessor systems?

A) Shared memory hierarchy within a chip
B) Presence of operating system
C) Use of interrupts
D) Availability of storage devices

21. Which sequence correctly describes process state transitions?

A) New → Ready → Running → Waiting → Ready → Terminated
B) Running → New → Ready → Waiting → Terminated
C) Waiting → New → Running → Ready → Terminated
D) Ready → New → Waiting → Running → Terminated

22. Which cloud deployment model is most appropriate for a government defense network requiring maximum control and isolation?

A) Public Cloud
B) Community Cloud
C) Private Cloud
D) Hybrid Cloud

23. In Hadoop MapReduce, which phase combines intermediate key-value pairs generated by mappers?

A) Input Split
B) Shuffle and Sort
C) Replication
D) Serialization

24. Assertion (A): Thrashing severely degrades system performance.
Reason (R): The system spends more time handling page faults than executing processes.

A) Both A and R are true and R explains A
B) Both A and R are true but R does not explain A
C) A is true but R is false
D) A is false but R is true

25. A disk scheduling graph shows the disk head moving only toward higher cylinder numbers before returning to the beginning. Which algorithm is represented?

A) FCFS
B) SSTF
C) SCAN
D) C–SCAN

26. A paging system uses a logical address space of 128 MB and a page size of 8 KB. How many bits are required for the page number field?

A) 10
B) 12
C) 14
D) 16

27. Which property of a monitor guarantees that only one process executes within the monitor at a given instant?

A) Reentrancy
B) Mutual Exclusion
C) Fragmentation
D) Address Binding

28. A cloud provider requires hardware-assisted virtualization with minimal overhead. Which component directly supports this requirement?

A) Compiler
B) Hypervisor–aware CPU Extension
C) Text Editor
D) Linker

29. Which combination correctly describes benefits of multithreading?

1. Improved responsiveness
2. Resource sharing
3. Increased communication overhead only
4. Scalability on multicore systems

A) 1 and 2 only
B) 1, 2 and 4 only
C) 2 and 3 only
D) 1, 3 and 4 only

30. Which expression correctly represents CPU Efficiency?

A) Useful CPU Time / Total CPU Time × 100
B) Waiting Time / Turnaround Time × 100
C) Response Time / Throughput
D) Throughput × Waiting Time

31. Which cloud security principle ensures that users receive only the minimum permissions necessary?

A) Redundancy Principle
B) Least Privilege Principle
C) Paging Principle
D) Hyperthreading Principle

32. All of the following are advantages of virtual memory EXCEPT:

A) Larger logical address space
B) Improved multiprogramming
C) Elimination of page faults
D) Better memory utilization

33. Cache memory access speed is commonly measured in:

A) Volts
B) Amperes
C) Nanoseconds
D) Hertz per byte

34. A banking system requires guaranteed exclusive access to account records during updates. Which synchronization mechanism is most suitable?

A) Mutex Lock
B) DMA Controller
C) Cache Line
D) Bootstrap Loader

35. Which category best describes Hadoop MapReduce?

A) Process Synchronization Tool
B) Distributed Data Processing Framework
C) Device Driver Framework
D) CPU Scheduling Algorithm

36. Which operating system structure organizes services into separate layers, where each layer depends only on lower layers?

A) Monolithic Structure
B) Layered Structure
C) Virtual Machine Structure
D) Hybrid Scheduling

37. In binary semaphores, the standard range of values is:

A) -1 and 1
B) 0 and 1
C) 1 and 2
D) Any Integer

38. A process flow diagram shows: Request Resource → Resource Unavailable → Wait Queue → Resource Granted → Execute. Which concept is represented?

A) Deadlock Recovery
B) Process Synchronization
C) Virtual Memory Allocation
D) File Allocation

39. A disk contains 200 cylinders numbered 0–199. The disk head is currently at cylinder 100 and receives a request for cylinder 150. What is the seek distance?

A) 25
B) 50
C) 75
D) 100

40. Which characteristic is most associated with a distributed file system?

A) Single–node dependency
B) Location Transparency
C) Sequential Scheduling
D) Interrupt Masking

41. Which storage virtualization component pools multiple physical storage devices into a single logical resource?

A) Hypervisor
B) Storage Controller
C) Compiler
D) Loader

42. Which are valid cloud service models?

1. IaaS
2. PaaS
3. SaaS
4. BIOSaaS

A) 1 and 2 only
B) 2 and 3 only
C) 1, 2 and 3 only
D) All four

43. Which formula determines average turnaround time?

A) Σ Turnaround Times / Number of Processes
B) Σ Waiting Times / Number of Processes
C) CPU Time / Throughput
D) Service Time / Burst Time

44. Which mechanism primarily protects against unauthorized access in cloud environments?

A) Authentication
B) Disk Defragmentation
C) Page Replacement
D) Interrupt Handling

45. Which is NOT a common page replacement algorithm?

A) FIFO
B) LRU
C) Optimal
D) FCFS

46. Network bandwidth in cloud environments is typically measured in:

A) Joules
B) Bits per Second
C) Newtons
D) Webers

47. A company wants to analyze petabytes of log data across thousands of servers. Which technology is most suitable?

A) Semaphore
B) Hadoop MapReduce
C) Mutex Lock
D) FCFS Scheduler

48. A hypervisor that runs directly on physical hardware is classified as:

A) Type-1 Hypervisor
B) Type-2 Hypervisor
C) Hybrid Hypervisor
D) Nested Loader

49. Which scheduling algorithm may suffer from starvation of long processes?

A) Round Robin
B) Priority Scheduling
C) FCFS
D) Multilevel Queue with Aging

50. In a typical 64-bit architecture, the size of a memory address is:

A) 16 bits
B) 32 bits
C) 64 bits
D) 128 bits

51. Which fundamental operating system principle ensures that multiple processes can safely share a resource without producing inconsistent results?

A) Locality of Reference
B) Mutual Exclusion
C) Virtualization
D) Fragmentation

52. Which memory allocation method divides memory into fixed-size blocks and logical memory into pages of equal size?

A) Segmentation
B) Paging
C) Swapping
D) Overlaying

53. A cloud data center integrates renewable energy to reduce operational costs. Which energy source is most commonly adopted for large-scale green data centers?

A) Diesel Generator
B) Coal-Based Plant
C) Solar Energy System
D) Internal Combustion Engine

54. Which property allows a cloud system to automatically increase computing resources during periods of high demand?

A) Portability
B) Elasticity
C) Fragmentation
D) Serialization

55. A process has a burst time of 40 ms. Under Round Robin scheduling with a quantum of 8 ms, how many CPU time slices are required?

A) 4
B) 5
C) 6
D) 8

56. Which statement correctly compares threads and processes?

A) Processes share all resources; threads do not
B) Threads have separate address spaces
C) Threads share process resources while processes have separate address spaces
D) Processes execute faster than threads in all cases

57. A distributed application exchanges messages without shared memory between nodes. Which IPC mechanism is most suitable?

A) Message Passing
B) Paging
C) Segmentation
D) DMA

58. Which page replacement algorithm replaces the page that has not been used for the longest time in the past?

A) FIFO
B) LRU
C) Optimal
D) SCAN

59. Which cloud service model provides virtual machines, storage, and networking resources directly to customers?

A) SaaS
B) PaaS
C) IaaS
D) DBaaS

60. Solid-state drives primarily use which memory technology?

A) Magnetic Core Memory
B) Flash Memory
C) Ferrite Memory
D) Vacuum Tube Storage

61. Statement I: Deadlock requires circular wait.
Statement II: Deadlock requires no preemption.
Statement III: Deadlock can occur without mutual exclusion.

A) I only
B) I and II only

C) II and III only
D) I, II and III

62. Which term describes excessive paging activity that significantly reduces system performance?

A) Starvation
B) Thrashing
C) Fragmentation
D) Multiprogramming

63. Cache hit ratio is defined as:

A) Misses / Total Accesses
B) Hits / Total Accesses
C) Access Time / Memory Size
D) Cache Size / Main Memory Size

64. Which principle explains why recently accessed data is likely to be accessed again soon?

A) Spatial Locality
B) Temporal Locality
C) Deadlock Avoidance
D) Virtualization

65. Which disk scheduling algorithm always selects the request closest to the current head position?

A) FCFS
B) SSTF
C) SCAN
D) C–SCAN

66. Which renewable energy source can provide continuous power generation even during nighttime operation of cloud data centers?

A) Solar Photovoltaic Only
B) Wind Energy
C) Coal Energy
D) Diesel Backup

67. Which cloud characteristic enables users to access services through standard internet-enabled devices?

A) Broad Network Access
B) Resource Starvation
C) Segmentation
D) Internal Fragmentation

68. A 64-bit architecture can theoretically address how many unique memory locations?

A) 2^{32}
B) 2^{48}
C) 2^{64}
D) 2^{128}

69. Which is the major advantage of segmentation over paging?

A) Eliminates page tables
B) Reflects logical program structure
C) Requires fixed-size memory blocks
D) Removes address translation

70. A company wants software applications delivered through a browser without local installation. Which service model should be selected?

A) IaaS
B) SaaS
C) PaaS
D) Virtualization Layer

71. Which scheduling algorithm is specifically designed to provide fairness among processes by allocating equal time quanta?

A) Priority Scheduling
B) Round Robin
C) SSTF
D) LRU

72. Which component in a virtualization environment directly manages guest operating systems?

A) Hypervisor
B) Compiler
C) Linker
D) Interpreter

73. Which memory type retains data even when power is removed?

A) SRAM
B) DRAM
C) Cache Memory
D) Flash Memory

74. Statement I: Virtualization improves resource utilization.
Statement II: Virtualization enables server consolidation.
Statement III: Virtualization eliminates the need for operating systems.

A) I only
B) I and II only
C) II and III only
D) I, II and III

75. The miss ratio of a cache is calculated as:

A) 1 − Hit Ratio
B) Hit Ratio × Access Time
C) Cache Size / Main Memory Size
D) Misses × Hits

UNIT V: OPERATING SYSTEMS AND CLOUD TECHNOLOGIES

ANSWER AND EXPLANATION

1. In a demand-paged virtual memory system, the Effective Access Time (EAT) is represented by a formula involving memory access time (m), page fault rate (p), and page fault service time (f). Which expression correctly represents the Effective Access Time?

A) EAT = (1 − p)m + pf

B) EAT = m + pf

C) EAT = p(m + f)

D) EAT = (m × f)/p

Answer: A) EAT = (1 − p)m + pf

Explanation: Effective Access Time measures the average memory access delay experienced by a process. The formula combines normal memory accesses and accesses that result in page faults. The term (1 − p)m represents successful memory references without page faults. The term pf accounts for the additional delay caused by servicing page faults.

2. A DMA controller transfers a block of 4096 bytes directly from an SSD to main memory over a 32-bit system bus. Assuming one bus transfer moves 32 bits of data and no wait states occur, how many bus transfer cycles are required to complete the DMA operation?

A) 512

B) 1024

C) 2048

D) 4096

Answer: B) 1024

Explanation: A 32-bit bus transfers 32 bits, which equals 4 bytes, per bus cycle. The total amount of data to be transferred is 4096 bytes. Therefore, the number of bus cycles required is 4096 ÷ 4 = 1024. DMA allows the transfer to occur without continuous CPU intervention, improving system efficiency.

3. A cloud data center is deploying hundreds of virtual machines on a multicore server. The administrator requires a software layer that can efficiently allocate CPU, memory, storage, and network resources among multiple guest operating systems while maintaining isolation. Which component is the most appropriate choice?

A) Device Driver

B) Hypervisor

C) Bootstrap Loader

D) Shell Interpreter

Answer: B) Hypervisor

Explanation: A hypervisor is specifically designed to create and manage virtual machines on physical hardware. It provides resource allocation, isolation, and scheduling among multiple guest operating systems. Device drivers only facilitate hardware communication and cannot manage virtualization environments. The bootstrap loader is involved only during system startup, while a shell interpreter provides a user interface for command execution. Therefore, the hypervisor is the correct component for large-scale virtualization deployment.

4. Which operating system component acts as an interface between user applications and hardware resources through controlled service requests?

A) Cache Controller
B) System Call Interface
C) DMA Controller
D) Bootstrap Program

Answer: B) System Call Interface

Explanation: The system call interface provides a controlled mechanism for user programs to request services from the operating system kernel. It ensures protection and resource management. User applications cannot directly access hardware in protected systems. The kernel validates requests before execution. Therefore, the system call interface serves as the gateway between applications and OS services.

5. Which of the following is NOT a characteristic objective of an operating system?

A) Resource allocation
B) Convenience for users
C) Program execution management
D) Increasing transistor density in processors

Answer: D) Increasing transistor density in processors

Explanation: Operating systems manage resources, execute programs, and provide convenient user environments. They do not influence semiconductor manufacturing processes. Transistor density is determined by hardware design and fabrication technologies. The OS operates above the hardware layer.

6. A system has a page size of 4 KB and logical address space of 64 MB. How many pages are present in the logical address space?

A) 8192
B) 16384
C) 32768
D) 65536

Answer: B) 16384

Explanation: Page size = 4 KB = 2^{12} bytes. Logical address space = 64 MB = 2^{26} bytes. Number of pages = 2^{26} / 2^{12} = 2^{14}. Thus the logical address space contains 16,384 pages. Paging divides memory into equal-sized blocks for efficient management.

7. Which property of semaphores allows synchronization among concurrent processes?

A) Mutual exclusion capability
B) Data compression capability
C) Address translation capability
D) Instruction decoding capability

Answer: A) Mutual exclusion capability

Explanation: Semaphores are synchronization primitives used to coordinate process execution. They help enforce mutual exclusion and prevent race conditions. Processes perform wait and signal operations on semaphores. This ensures orderly access to shared resources. Therefore mutual exclusion is their primary synchronization property.

8. Identify the correct sequence during interrupt processing.

A) Save Context → Execute ISR → Restore Context → Resume Process
B) Execute ISR → Save Context → Resume Process → Restore Context
C) Restore Context → Execute ISR → Save Context → Resume Process
D) Resume Process → Save Context → Execute ISR → Restore Context

Answer: A) Save Context → Execute ISR → Restore Context → Resume Process

Explanation: When an interrupt occurs, the processor first saves the current execution state. The interrupt service routine (ISR) is then executed. After completion, the saved context is restored. Finally, execution resumes from the interrupted process. This sequence guarantees correct program continuity.

9. A real-time industrial controller requires deterministic scheduling and bounded response times. Which scheduling algorithm is most suitable?

A) First Come First Served
B) Round Robin
C) Earliest Deadline First
D) Random Scheduling

Answer: C) Earliest Deadline First

Explanation: Earliest Deadline First dynamically prioritizes tasks with the closest deadlines. It is widely used in real-time systems requiring predictable responses. FCFS and Round Robin focus mainly on fairness rather than deadline guarantees. Random scheduling provides no predictability. Therefore EDF is best suited for deterministic environments.

10. Which term refers to a process currently being executed by the CPU?

A) Ready Process
B) Blocked Process
C) Running Process
D) Zombie Process

Answer: C) Running Process

Explanation: A running process is actively executing instructions on the CPU. Ready processes are waiting for CPU allocation. Blocked processes wait for events such as I/O completion. Zombie processes have terminated but still retain process table entries. Hence the running state corresponds to active execution.

11. Assertion (A): Multithreading improves responsiveness in applications.
Reason (R): Multiple threads can execute concurrently while sharing process resources.

A) Both A and R are true and R explains A
B) Both A and R are true but R does not explain A
C) A is true but R is false
D) A is false but R is true

Answer: A) Both A and R are true and R explains A

Explanation: Multithreading allows separate execution flows within a process. Threads share resources while performing tasks independently. This enables background operations without blocking user interactions. As a result, applications remain responsive. Therefore the reason correctly explains the assertion.

12. A CPU utilization graph shows utilization increasing from 30% to 90% after introducing multiprogramming. What is the most likely reason?

A) Increased context switching only
B) Better overlap of CPU and I/O activities
C) Reduced memory size
D) Increased interrupt latency

Answer: B) Better overlap of CPU and I/O activities

Explanation: Multiprogramming allows another process to execute when one waits for I/O. This minimizes CPU idle time. Consequently overall utilization increases significantly. Reduced memory and higher interrupt latency would generally degrade performance. Thus overlapping CPU and I/O operations is the main reason.

13. A Round Robin scheduler uses a time quantum of 5 ms. A process requires 23 ms CPU time. Assuming no overhead, how many CPU allocations will the process receive?

A) 3
B) 4
C) 5
D) 6

Answer: C) 5

Explanation: The process requires 23 ms. Each allocation provides 5 ms except the last one. The allocations are 5 + 5 + 5 + 5 + 3 = 23 ms. Therefore five CPU allocations are needed. Round Robin ensures fairness by cyclic time sharing.

14. Which formula represents CPU utilization?

A) Busy Time / Total Time × 100
B) Idle Time / Busy Time × 100
C) Turnaround Time / Waiting Time × 100
D) Throughput / Response Time

Answer: A) Busy Time / Total Time × 100

Explanation: CPU utilization measures the percentage of time the processor remains busy. It is calculated by dividing busy time by total observation time. Higher utilization generally indicates effective resource use. The remaining formulas represent unrelated performance metrics.

15. A cache memory has an access time of 5 ns and main memory has 80 ns access time. If cache hit ratio is 90%, find average memory access time.

A) 12.5 ns
B) 15 ns
C) 20 ns
D) 25 ns

Answer: A) 12.5 ns

Explanation: Average access time = (0.9 × 5) + (0.1 × 80). This equals 4.5 + 8 = 12.5 ns. A high hit ratio significantly improves memory performance. Cache memory reduces access delays for frequently used data. Therefore the average access time is 12.5 ns.

16. Which storage technology is most suitable for implementing a high-speed paging device in cloud servers?

A) Magnetic Tape
B) Optical Disk
C) NVMe SSD
D) Punch Card

Answer: C) NVMe SSD

Explanation: Paging devices require extremely low access latency and high throughput. NVMe SSDs provide fast random access operations compared to traditional storage. Magnetic tapes and optical disks are unsuitable for intensive paging workloads. Punch cards are obsolete. Therefore NVMe SSDs are the best choice.

17. What is the primary role of a monitor in process synchronization?

A) Memory allocation
B) Encapsulating shared data and synchronization operations
C) Disk scheduling
D) Address translation

Answer: B) Encapsulating shared data and synchronization operations

Explanation: Monitors provide a high-level synchronization mechanism. Shared variables and synchronization procedures are grouped together. Only one process can execute a monitor procedure at a time. This simplifies concurrent programming. Hence monitors help manage shared resource access safely.

18. Which of the following is NOT a deadlock necessary condition?

A) Mutual Exclusion
B) Hold and Wait
C) Circular Wait
D) Priority Inversion

Answer: D) Priority Inversion

Explanation: The four necessary deadlock conditions are mutual exclusion, hold-and-wait, no preemption, and circular wait. Priority inversion is a scheduling problem. It may affect responsiveness but is not a deadlock condition. Deadlocks require all four necessary conditions simultaneously. Therefore priority inversion is the exception.

19. A 32-bit system uses byte addressing. What is the maximum directly addressable memory space?

A) 2 GB
B) 4 GB
C) 8 GB
D) 16 GB

Answer: B) 4 GB

Explanation: A 32-bit address can represent 2^{32} unique addresses. Since each address refers to one byte, the total addressable memory equals 2^{32} bytes. This corresponds to 4 GB. This limitation motivated the adoption of 64-bit architectures.

20.Which property distinguishes multicore processors from multiprocessor systems?

A) Shared memory hierarchy within a chip
B) Presence of operating system
C) Use of interrupts
D) Availability of storage devices

Answer: A) Shared memory hierarchy within a chip

Explanation: Multicore processors integrate multiple cores on a single chip. These cores often share cache and memory structures. Multiprocessor systems may contain separate physical processors. The remaining options are common to both architectures. Therefore shared on-chip resources distinguish multicore systems.

21. Which sequence correctly describes process state transitions?

A) New → Ready → Running → Waiting → Ready → Terminated
B) Running → New → Ready → Waiting → Terminated
C) Waiting → New → Running → Ready → Terminated
D) Ready → New → Waiting → Running → Terminated

Answer: A) New → Ready → Running → Waiting → Ready → Terminated

Explanation: A process is created in the new state. It enters the ready queue and waits for CPU allocation. After execution it may move to waiting for I/O events. Upon event completion it returns to ready. Finally it terminates after completing execution.

22. Which cloud deployment model is most appropriate for a government defense network requiring maximum control and isolation?

A) Public Cloud
B) Community Cloud
C) Private Cloud
D) Hybrid Cloud

Answer: C) Private Cloud

Explanation: Private clouds provide dedicated infrastructure and greater administrative control. Sensitive government applications require strict security and isolation. Public clouds involve shared environments. Hybrid and community models may not provide the same level of exclusive control. Therefore a private cloud is most suitable.

23. In Hadoop MapReduce, which phase combines intermediate key–value pairs generated by mappers?

A) Input Split
B) Shuffle and Sort
C) Replication
D) Serialization

Answer: B) Shuffle and Sort

Explanation: After mapping, intermediate data must be grouped according to keys. The shuffle and sort phase performs this organization. Reducers then process grouped values. This stage is fundamental to distributed computation.

24. Assertion (A): Thrashing severely degrades system performance.
Reason (R): The system spends more time handling page faults than executing processes.

A) Both A and R are true and R explains A
B) Both A and R are true but R does not explain A
C) A is true but R is false
D) A is false but R is true

Answer: A) Both A and R are true and R explains A

Explanation: Thrashing occurs when insufficient memory causes excessive paging activity. The processor repeatedly services page faults. Useful program execution decreases drastically. System throughput falls sharply. Therefore the reason accurately explains the assertion.

25. A disk scheduling graph shows the disk head moving only toward higher cylinder numbers before returning to the beginning. Which algorithm is represented?

A) FCFS
B) SSTF
C) SCAN
D) C–SCAN

Answer: D) C–SCAN

Explanation: C–SCAN services requests in a single direction. When the highest cylinder is reached, the head returns to the beginning without servicing requests during the return. This provides more uniform waiting times. SCAN services requests in both directions. Therefore the graph represents the C–SCAN algorithm.

26. A paging system uses a logical address space of 128 MB and a page size of 8 KB. How many bits are required for the page number field?

A) 10
B) 12
C) 14
D) 16

Answer: C) 14

Explanation: The page size is 8 KB = 2^{13} bytes. The logical address space is 128 MB = 2^{27} bytes. Therefore, the number of pages is $2^{27}/2^{13} = 2^{14}$. Hence, 14 bits are needed to represent the page number. This calculation is fundamental in paging–based memory management systems.

27. Which property of a monitor guarantees that only one process executes within the monitor at a given instant?

A) Reentrancy
B) Mutual Exclusion
C) Fragmentation
D) Address Binding

Answer: B) Mutual Exclusion

Explanation: Monitors are high–level synchronization constructs. They inherently enforce mutual exclusion on shared resources. Only one thread or process may execute monitor procedures simultaneously. This prevents race conditions. Therefore, mutual exclusion is the key theoretical property.

28. A cloud provider requires hardware-assisted virtualization with minimal overhead. Which component directly supports this requirement?

A) Compiler
B) Hypervisor-aware CPU Extension
C) Text Editor
D) Linker

Answer: B) Hypervisor-aware CPU Extension

Explanation: Modern processors include virtualization extensions such as Intel VT-x and AMD-V. These allow efficient execution of guest operating systems. Hypervisors utilize these extensions to reduce virtualization overhead. Traditional software tools such as compilers and linkers are unrelated. Thus CPU virtualization extensions are required.

29. Which combination correctly describes benefits of multithreading?

1. Improved responsiveness
2. Resource sharing
3. Increased communication overhead only
4. Scalability on multicore systems

A) 1 and 2 only
B) 1, 2 and 4 only
C) 2 and 3 only
D) 1, 3 and 4 only

Answer: B) 1, 2 and 4 only

Explanation: Multithreading improves responsiveness by allowing concurrent execution. Threads share process resources efficiently. Multicore systems benefit from thread-level parallelism. Increased communication overhead only is not considered a benefit. Therefore statements 1, 2, and 4 are correct.

30. Which expression correctly represents CPU Efficiency?

A) Useful CPU Time / Total CPU Time × 100
B) Waiting Time / Turnaround Time × 100
C) Response Time / Throughput
D) Throughput × Waiting Time

Answer: A) Useful CPU Time / Total CPU Time × 100

Explanation: CPU efficiency measures the percentage of productive processing time. It compares useful execution time against total available CPU time. High efficiency indicates minimal idle or wasted cycles. The other formulas describe unrelated performance measures.

31. Which cloud security principle ensures that users receive only the minimum permissions necessary?

A) Redundancy Principle
B) Least Privilege Principle
C) Paging Principle
D) Hyperthreading Principle

Answer: B) Least Privilege Principle

Explanation: Least privilege minimizes security risks by restricting permissions. Users and applications receive only the access required for their tasks. This reduces attack surfaces and accidental misuse. It is a fundamental cybersecurity principle.

32. All of the following are advantages of virtual memory EXCEPT:

A) Larger logical address space
B) Improved multiprogramming
C) Elimination of page faults
D) Better memory utilization

Answer: C) Elimination of page faults

Explanation: Virtual memory enables execution of programs larger than physical memory. It improves memory utilization and multiprogramming levels. However, page faults are inherent to demand paging systems. Virtual memory manages page faults rather than eliminating them.

33. Cache memory access speed is commonly measured in:

A) Volts
B) Amperes
C) Nanoseconds
D) Hertz per byte

Answer: C) Nanoseconds

Explanation: Memory access time indicates how quickly data can be retrieved. Cache memories typically operate within nanosecond ranges. Volts and amperes measure electrical quantities. Hertz per byte is not a standard memory timing unit. Therefore nanoseconds are used.

34. A banking system requires guaranteed exclusive access to account records during updates. Which synchronization mechanism is most suitable?

A) Mutex Lock
B) DMA Controller
C) Cache Line
D) Bootstrap Loader

Answer: A) Mutex Lock

Explanation: Mutex locks provide exclusive access to shared resources. Only one thread can hold a mutex at a time. This prevents simultaneous updates that could corrupt account balances. DMA and cache mechanisms do not provide synchronization. Hence mutex locks are appropriate.

35. Which category best describes Hadoop MapReduce?

A) Process Synchronization Tool
B) Distributed Data Processing Framework
C) Device Driver Framework
D) CPU Scheduling Algorithm

Answer: B) Distributed Data Processing Framework

Explanation: Hadoop MapReduce is designed for large-scale distributed computation. It processes massive datasets across clusters of machines. It is not a synchronization tool or scheduling algorithm. The framework divides tasks into map and reduce phases. Thus it is classified as a distributed processing framework.

36. Which operating system structure organizes services into separate layers, where each layer depends only on lower layers?

A) Monolithic Structure
B) Layered Structure

C) Virtual Machine Structure
D) Hybrid Scheduling

Answer: B) Layered Structure

Explanation: Layered operating systems divide functionality into hierarchical levels. Each layer interacts only with adjacent layers. This improves modularity and maintainability. Monolithic systems do not enforce strict layer separation.

37. In binary semaphores, the standard range of values is:

A) -1 and 1
B) 0 and 1
C) 1 and 2
D) Any Integer

Answer: B) 0 and 1

Explanation: Binary semaphores function similarly to locks. Their values are restricted to 0 and 1. A value of 1 generally indicates availability. A value of 0 indicates resource occupancy. Hence the standard range is 0 and 1.

38. A process flow diagram shows: Request Resource → Resource Unavailable → Wait Queue → Resource Granted → Execute. Which concept is represented?

A) Deadlock Recovery
B) Process Synchronization
C) Virtual Memory Allocation
D) File Allocation

Answer: B) Process Synchronization

Explanation: The flow illustrates controlled access to shared resources. Processes wait when resources are unavailable. Once granted, execution proceeds safely. Such coordination is the goal of synchronization mechanisms. Therefore the diagram represents process synchronization.

39. A disk contains 200 cylinders numbered 0–199. The disk head is currently at cylinder 100 and receives a request for cylinder 150. What is the seek distance?

A) 25
B) 50
C) 75
D) 100

Answer: B) 50

Explanation: Seek distance equals the absolute difference between the current and requested cylinder. Here, |150 − 100| = 50. Disk scheduling algorithms often use this value. Lower seek distances generally improve performance..

40. Which characteristic is most associated with a distributed file system?

A) Single-node dependency
B) Location Transparency
C) Sequential Scheduling
D) Interrupt Masking

Answer: B) Location Transparency

Explanation: Distributed file systems hide physical storage locations from users. Files appear accessible regardless of actual storage nodes. This feature is called location transparency. It simplifies resource access in distributed environments.

41. Which storage virtualization component pools multiple physical storage devices into a single logical resource?

A) Hypervisor
B) Storage Controller
C) Compiler
D) Loader

Answer: B) Storage Controller

Explanation: Storage virtualization aggregates multiple storage devices. A storage controller manages the abstraction layer. Users interact with logical storage rather than individual devices. This improves scalability and flexibility. Therefore storage controllers enable pooling.

42. Which are valid cloud service models?

1. IaaS
2. PaaS
3. SaaS
4. BIOSaaS

A) 1 and 2 only
B) 2 and 3 only
C) 1, 2 and 3 only
D) All four

Answer: C) 1, 2 and 3 only

Explanation: IaaS, PaaS, and SaaS are recognized cloud service models. They provide infrastructure, platforms, and software respectively. BIOSaaS is not a standard cloud model. Cloud architectures are commonly categorized into these three service layers.

43. Which formula determines average turnaround time?

A) Σ Turnaround Times / Number of Processes
B) Σ Waiting Times / Number of Processes
C) CPU Time / Throughput
D) Service Time / Burst Time

Answer: A) Σ Turnaround Times / Number of Processes

Explanation: Turnaround time measures total completion duration. Average turnaround time is obtained by summing all turnaround times and dividing by the number of processes. It is a key scheduling performance metric. Lower values indicate improved efficiency.

44. Which mechanism primarily protects against unauthorized access in cloud environments?

A) Authentication
B) Disk Defragmentation
C) Page Replacement
D) Interrupt Handling

Answer: A) Authentication

Explanation: Authentication verifies user identity before granting access. It is the first line of defense in cloud security. Without authentication, unauthorized access becomes

possible. Disk management and paging mechanisms do not provide identity verification. Thus authentication is essential.

45. Which is NOT a common page replacement algorithm?

A) FIFO
B) LRU
C) Optimal
D) FCFS

Answer: D) FCFS

Explanation: FIFO, LRU, and Optimal are standard page replacement algorithms. They determine which memory page should be removed. FCFS is a CPU scheduling algorithm rather than a page replacement strategy. Therefore FCFS is the exception.

46. Network bandwidth in cloud environments is typically measured in:

A) Joules
B) Bits per Second
C) Newtons
D) Webers

Answer: B) Bits per Second

Explanation: Bandwidth indicates the rate of data transfer. The standard unit is bits per second (bps). Larger networks use Mbps, Gbps, or Tbps. Joules, Newtons, and Webers measure unrelated physical quantities. Hence bits per second is correct.

47. A company wants to analyze petabytes of log data across thousands of servers. Which technology is most suitable?

A) Semaphore
B) Hadoop MapReduce
C) Mutex Lock
D) FCFS Scheduler

Answer: B) Hadoop MapReduce

Explanation: MapReduce is designed for distributed processing of massive datasets. It divides work among many nodes and aggregates results efficiently. Semaphores and

mutexes are synchronization tools. FCFS is a scheduling policy. Therefore Hadoop MapReduce is the best choice.

48. A hypervisor that runs directly on physical hardware is classified as:

A) Type–1 Hypervisor
B) Type–2 Hypervisor
C) Hybrid Hypervisor
D) Nested Loader

Answer: A) Type–1 Hypervisor

Explanation: Type–1 hypervisors operate directly on hardware. They provide high performance and strong isolation. Examples include enterprise virtualization platforms. Type–2 hypervisors run on top of host operating systems. Therefore the correct classification is Type–1.

49. Which scheduling algorithm may suffer from starvation of long processes?

A) Round Robin
B) Priority Scheduling
C) FCFS
D) Multilevel Queue with Aging

Answer: B) Priority Scheduling

Explanation: Priority scheduling favors higher–priority processes. Lower–priority processes may wait indefinitely if higher–priority tasks continue arriving. This phenomenon is known as starvation. Aging is often introduced to mitigate the problem. Hence priority scheduling is the answer.

50. In a typical 64–bit architecture, the size of a memory address is:

A) 16 bits
B) 32 bits
C) 64 bits
D) 128 bits

Answer: C) 64 bits

Explanation: A 64-bit architecture is designed to process and represent 64-bit addresses. This significantly expands the addressable memory space compared to 32-bit systems. Modern servers and cloud platforms commonly use 64-bit architectures. Therefore the address size is 64 bits.

51. Which fundamental operating system principle ensures that multiple processes can safely share a resource without producing inconsistent results?

A) Locality of Reference
B) Mutual Exclusion
C) Virtualization
D) Fragmentation

Answer: B) Mutual Exclusion

Explanation: Mutual exclusion is a fundamental synchronization principle in operating systems. It ensures that only one process accesses a critical resource at a time. Without it, race conditions may occur, leading to inconsistent data. Synchronization primitives such as mutexes and semaphores implement this concept.

52. Which memory allocation method divides memory into fixed-size blocks and logical memory into pages of equal size?

A) Segmentation
B) Paging
C) Swapping
D) Overlaying

Answer: B) Paging

Explanation: Paging divides physical memory into frames and logical memory into pages. Each page can be loaded into any available frame. This eliminates external fragmentation. Address translation is performed using page tables.

53. A cloud data center integrates renewable energy to reduce operational costs. Which energy source is most commonly adopted for large-scale green data centers?

A) Diesel Generator
B) Coal-Based Plant

C) Solar Energy System
D) Internal Combustion Engine

Answer: C) Solar Energy System

Explanation: Many modern cloud providers use solar energy to power data centers. Solar systems reduce dependence on fossil fuels and lower carbon emissions. Renewable energy also supports sustainability goals. Coal and diesel are non-renewable and environmentally harmful. Therefore solar energy is the preferred choice.

54. Which property allows a cloud system to automatically increase computing resources during periods of high demand?

A) Portability
B) Elasticity
C) Fragmentation
D) Serialization

Answer: B) Elasticity

Explanation: Elasticity is a defining property of cloud computing. It enables dynamic scaling of resources according to workload demands. During peak usage, additional resources can be provisioned automatically. This improves performance and cost efficiency.

55. A process has a burst time of 40 ms. Under Round Robin scheduling with a quantum of 8 ms, how many CPU time slices are required?

A) 4
B) 5
C) 6
D) 8

Answer: B) 5

Explanation: Each time slice provides 8 ms of CPU execution. The process requires 40 ms total. Dividing 40 by 8 gives 5 complete quanta. No additional partial quantum is needed.

56. Which statement correctly compares threads and processes?

A) Processes share all resources; threads do not
B) Threads have separate address spaces
C) Threads share process resources while processes have separate address spaces
D) Processes execute faster than threads in all cases

Answer: C) Threads share process resources while processes have separate address spaces

Explanation: Threads within the same process share memory and resources. Processes generally maintain independent address spaces. Thread creation and switching are typically less expensive than process management. Therefore option C accurately compares them.

57. A distributed application exchanges messages without shared memory between nodes. Which IPC mechanism is most suitable?

A) Message Passing
B) Paging
C) Segmentation
D) DMA

Answer: A) Message Passing

Explanation: Message passing is widely used in distributed systems. It enables processes on separate machines to communicate safely. Shared memory is not required. The method provides synchronization and data exchange capabilities. Thus message passing is the best choice.

58. Which page replacement algorithm replaces the page that has not been used for the longest time in the past?

A) FIFO
B) LRU
C) Optimal
D) SCAN

Answer: B) LRU

Explanation: Least Recently Used (LRU) tracks past memory references. The page unused for the longest duration is selected for replacement. LRU approximates the principle of locality. It generally performs better than FIFO.

59. Which cloud service model provides virtual machines, storage, and networking resources directly to customers?

A) SaaS
B) PaaS
C) IaaS
D) DBaaS

Answer: C) IaaS

Explanation: Infrastructure as a Service delivers virtualized computing resources. Users can create virtual machines and configure networking and storage. It provides maximum flexibility among cloud service models. SaaS and PaaS operate at higher abstraction levels.

60. Solid-state drives primarily use which memory technology?

A) Magnetic Core Memory
B) Flash Memory
C) Ferrite Memory
D) Vacuum Tube Storage

Answer: B) Flash Memory

Explanation: SSDs store data using non-volatile flash memory cells. These devices provide high speed and low latency. Unlike hard disks, SSDs contain no moving parts. This improves reliability and performance.

61. Statement I: Deadlock requires circular wait.
Statement II: Deadlock requires no preemption.
Statement III: Deadlock can occur without mutual exclusion.

A) I only
B) I and II only
C) II and III only
D) I, II and III

Answer: B) I and II only

Explanation: Circular wait and no preemption are necessary deadlock conditions. Mutual exclusion is also required for deadlock. Therefore Statement III is false. Deadlock cannot occur if mutual exclusion is absent. Hence Statements I and II alone are correct.

62. Which term describes excessive paging activity that significantly reduces system performance?

A) Starvation
B) Thrashing
C) Fragmentation
D) Multiprogramming

Answer: B) Thrashing

Explanation: Thrashing occurs when the system spends most of its time servicing page faults. Very little useful work is completed. It is often caused by insufficient physical memory. CPU utilization may decrease sharply.

63. Cache hit ratio is defined as:

A) Misses / Total Accesses
B) Hits / Total Accesses
C) Access Time / Memory Size
D) Cache Size / Main Memory Size

Answer: B) Hits / Total Accesses

Explanation: The hit ratio measures cache effectiveness. It represents the proportion of memory requests satisfied directly from cache. Higher hit ratios improve performance significantly. Miss ratio is the complement of hit ratio.

64. Which principle explains why recently accessed data is likely to be accessed again soon?

A) Spatial Locality
B) Temporal Locality
C) Deadlock Avoidance
D) Virtualization

Answer: B) Temporal Locality

Explanation: Temporal locality is a fundamental memory access principle. Data referenced recently tends to be reused shortly afterward. Cache systems exploit this behavior. It improves memory performance and reduces access latency.

65. Which disk scheduling algorithm always selects the request closest to the current head position?

A) FCFS
B) SSTF
C) SCAN
D) C–SCAN

Answer: B) SSTF

Explanation: Shortest Seek Time First minimizes immediate seek distance. The algorithm chooses the nearest pending request. This often reduces average seek time. However, it may cause starvation for distant requests.

66. Which renewable energy source can provide continuous power generation even during nighttime operation of cloud data centers?

A) Solar Photovoltaic Only
B) Wind Energy
C) Coal Energy
D) Diesel Backup

Answer: B) Wind Energy

Explanation: Wind turbines can generate electricity whenever sufficient wind is available. Unlike solar systems, generation is not limited to daylight hours. Large cloud providers frequently combine wind and solar resources. This supports sustainable operations.

67. Which cloud characteristic enables users to access services through standard internet-enabled devices?

A) Broad Network Access
B) Resource Starvation
C) Segmentation
D) Internal Fragmentation

Answer: A) Broad Network Access

Explanation: Broad network access is one of the essential cloud characteristics. Resources are accessible over networks using standard protocols. Users can connect from laptops, tablets, and smartphones. This enhances flexibility and accessibility.

68. A 64-bit architecture can theoretically address how many unique memory locations?

A) 2^{32}
B) 2^{48}
C) 2^{64}
D) 2^{128}

Answer: C) 2^{64}

Explanation: A 64-bit address contains 64 binary digits. Each bit can represent two states. Therefore the total number of unique addresses is 2^{64}. This allows an extremely large address space.

69. Which is the major advantage of segmentation over paging?

A) Eliminates page tables
B) Reflects logical program structure
C) Requires fixed-size memory blocks
D) Removes address translation

Answer: B) Reflects logical program structure

Explanation: Segmentation divides memory according to logical units such as functions or modules. This matches the programmer's perspective. Paging uses fixed-size blocks without semantic meaning. Therefore segmentation provides a more natural program representation.

70. A company wants software applications delivered through a browser without local installation. Which service model should be selected?

A) IaaS
B) SaaS
C) PaaS
D) Virtualization Layer

Answer: B) SaaS

Explanation: Software as a Service provides complete applications over the internet. Users access services through web browsers. Installation and maintenance are handled by the provider. Examples include cloud-based office applications.

71. Which scheduling algorithm is specifically designed to provide fairness among processes by allocating equal time quanta?

A) Priority Scheduling
B) Round Robin
C) SSTF
D) LRU

Answer: B) Round Robin

Explanation: Round Robin scheduling distributes CPU time equally among ready processes. Each process receives a fixed time quantum. This approach prevents monopolization of CPU resources. It is commonly used in time-sharing systems.

72. Which component in a virtualization environment directly manages guest operating systems?

A) Hypervisor
B) Compiler
C) Linker
D) Interpreter

Answer: A) Hypervisor

Explanation: The hypervisor creates and controls virtual machines. It allocates CPU, memory, storage, and networking resources. Guest operating systems execute within the virtualized environment. This is the core component of virtualization platforms. Therefore the answer is hypervisor.

73. Which memory type retains data even when power is removed?

A) SRAM
B) DRAM
C) Cache Memory
D) Flash Memory

Answer: D) Flash Memory

Explanation: Flash memory is non–volatile. Stored information remains available after power loss. SRAM and DRAM are volatile memory technologies. They lose contents when power is disconnected.

74. Statement I: Virtualization improves resource utilization.
Statement II: Virtualization enables server consolidation.
Statement III: Virtualization eliminates the need for operating systems.

A) I only
B) I and II only
C) II and III only
D) I, II and III

Answer: B) I and II only

Explanation: Virtualization allows multiple workloads to share hardware efficiently. It supports server consolidation and improved utilization. However, guest operating systems are still required. Therefore Statement III is incorrect.

75. The miss ratio of a cache is calculated as:

A) 1 − Hit Ratio
B) Hit Ratio × Access Time
C) Cache Size / Main Memory Size
D) Misses × Hits

Answer: A) 1 − Hit Ratio

Explanation: Cache accesses result in either hits or misses. The sum of hit ratio and miss ratio equals one. Therefore miss ratio is obtained by subtracting hit ratio from one. This metric helps evaluate cache performance.

76. In a demand paging system, if page fault probability is p, memory access time is ma, and page fault service time is pf, which formula correctly represents Effective Access Time (EAT)?

A) EAT = (1 − p) × ma + p × pf
B) EAT = ma × pf
C) EAT = ma + pf
D) EAT = pf / ma