Final Exam
Term 251

Course title: Distributed Systems **Course Code**: CS 435

Exam date: 13 December 2025 **Exam Time: 120** minutes

Student Name: _____

Student ID:									

Circle the Section time: **10 AM** **11 AM**

Question No.	**Part a**	**Part b**	**Part c**	**Student's Score**
Question 1 (CLO 1)	/4	/4		**/8**
Question 2 (CLO 2)	/3	/6		**/9**
Question 3 (CLO 3)	/3	/4	/5	**/12**
Question 4 (CLO 4)	/5			**/5**
Question 5 (CLO 5)	/6			**/6**
Total				**/40**

Instructions:

- This exam contains five questions with multiple parts. There are 7 sheets in total, each sheet has two sides.
- **Time allowed: 120 minutes**
- Closed Book, Closed Notes.
- Use of computing devices / smartphones etc is strictly prohibited. Students can use calculators.
- Answer the problems on the exam sheets only. No additional attachments would be accepted.
- If you need extra space use the last page/sheet of this set of papers.

This page is left blank intentionally

Please record your answers to all multiple-choice questions only in the designated answer bubbles below. Answers written anywhere else will not be graded.

Each MCQ in the box is weighted 1 point. Otherwise it is 0.5 weighted points.

Question 1

1. Ⓐ Ⓑ Ⓒ Ⓓ
2. Ⓐ Ⓑ Ⓒ Ⓓ
3. Ⓐ Ⓑ Ⓒ Ⓓ
4. Ⓐ Ⓑ Ⓒ Ⓓ
5. Ⓐ Ⓑ Ⓒ Ⓓ
6. Ⓐ Ⓑ Ⓒ Ⓓ
7. Ⓐ Ⓑ Ⓒ Ⓓ
8. Ⓐ Ⓑ Ⓒ Ⓓ
9. Ⓐ Ⓑ Ⓒ Ⓓ
10. Ⓐ Ⓑ Ⓒ Ⓓ
11. Ⓐ Ⓑ Ⓒ Ⓓ
12. Ⓐ Ⓑ Ⓒ Ⓓ

Question 2

13. Ⓐ Ⓑ Ⓒ Ⓓ
14. Ⓐ Ⓑ Ⓒ Ⓓ
15. Ⓐ Ⓑ Ⓒ Ⓓ
16. Ⓐ Ⓑ Ⓒ Ⓓ
17. Ⓐ Ⓑ Ⓒ Ⓓ
18. Ⓐ Ⓑ Ⓒ Ⓓ
19. Ⓐ Ⓑ Ⓒ Ⓓ
20. Ⓐ Ⓑ Ⓒ Ⓓ
21. Ⓐ Ⓑ Ⓒ Ⓓ
22. Ⓐ Ⓑ Ⓒ Ⓓ
23. Ⓐ Ⓑ Ⓒ Ⓓ
24. Ⓐ Ⓑ Ⓒ Ⓓ

Question 3

25. Ⓐ Ⓑ Ⓒ Ⓓ
26. Ⓐ Ⓑ Ⓒ Ⓓ
27. Ⓐ Ⓑ Ⓒ Ⓓ
28. Ⓐ Ⓑ Ⓒ Ⓓ
29. Ⓐ Ⓑ Ⓒ Ⓓ
30. Ⓐ Ⓑ Ⓒ Ⓓ
31. Ⓐ Ⓑ Ⓒ Ⓓ
32. Ⓐ Ⓑ Ⓒ Ⓓ
33. Ⓐ Ⓑ Ⓒ Ⓓ
34. Ⓐ Ⓑ Ⓒ Ⓓ
35. Ⓐ Ⓑ Ⓒ Ⓓ
36. Ⓐ Ⓑ Ⓒ Ⓓ
37. Ⓐ Ⓑ Ⓒ Ⓓ
38. Ⓐ Ⓑ Ⓒ Ⓓ
39. Ⓐ Ⓑ Ⓒ Ⓓ

Question 4

40. Ⓐ Ⓑ Ⓒ Ⓓ
41. Ⓐ Ⓑ Ⓒ Ⓓ
42. Ⓐ Ⓑ Ⓒ Ⓓ
43. Ⓐ Ⓑ Ⓒ Ⓓ
44. Ⓐ Ⓑ Ⓒ Ⓓ

Question 5

45. Ⓐ Ⓑ Ⓒ Ⓓ
46. Ⓐ Ⓑ Ⓒ Ⓓ
47. Ⓐ Ⓑ Ⓒ Ⓓ
48. Ⓐ Ⓑ Ⓒ Ⓓ
49. Ⓐ Ⓑ Ⓒ Ⓓ
50. Ⓐ Ⓑ Ⓒ Ⓓ

.

Question 1 [8 points - CLO 1]

Part A [4]- Select the most appropriate choice from the following MCQs

1	Which of the following is a primary goal of a distributed system? A) Low Latency and High Bandwidth B) Scalability, Availability, and Fault Tolerance C) Centralized Control and Simplified Management D) Homogeneous Hardware and Software
2	Which of the following is an example of a structured overlay network that uses a DHT? A) BitTorrent's unstructured peer discovery B) Chord C) Wi-Fi routing D) Ethernet switching
3	Which statement accurately captures the difference between a Distributed Operating System (DOS) and a Network Operating System (NOS)? A) A DOS requires users to explicitly mount remote file systems, while a NOS makes remote files appear local automatically. B) A DOS only works on heterogeneous hardware, while a NOS requires homogeneous nodes. C) A DOS provides a single, coherent system view by managing resources globally, while a NOS manages independent nodes that share resources explicitly. D) A DOS has higher overhead and is less scalable than a NOS because it uses a microkernel design.
4	Which of the following is the main reason to choose an event-based architecture for a real-time stock trading dashboard? A) To let different parts of the dashboard get instant updates as soon as something happens, without tightly connecting them. B) To keep all past trading data in one central place. C) To make sure money-related updates (like account balances) are handled with strict accuracy. D) To force all data to follow one fixed, controlled path for auditing.
5	A well know problem in computer architecture is the "Power Wall", what is it? A) Increasing clock speeds leads to unsustainable heat and energy consumption. B) Data centers does not have enough electrical power to run. C) Mobile devices have limited battery life. D) Processors require more power than memory modules.
6	The Fork/Join framework uses a work-stealing algorithm primarily to: A) Reduce memory consumption. B) Ensure tasks are executed in the order they were created. C) Improve load balancing by having idle threads steal work from busy ones. D) Prevent race conditions on shared variables.
7	What is a key characteristic of a RESTful web service? A) It uses complex, binary protocols for efficiency. B) It is built on top of the HTTP protocol and its methods (GET, POST, etc.). C) It requires a WSDL file for service description. D) It only works with XML data format.
8	A hybrid P2P/Client-Server architecture is often praised for balancing the strengths of both models. However, this combination introduces a unique, inherent challenge. What is it? A) It sacrifices the simplicity of the P2P model and Inherits complexity of the single point of failure. B) It forces all peers to use a homogeneous operating system, negating the hardware flexibility of P2P. C) It cannot provide any form of location or access transparency to the end-user. D) It inherits the single point of failure from the client-server model and the coordination complexity from the P2P model, leading to increased system complexity and overhead.

Part B [4]- Read the following scenario and then select the most appropriate choice from the following MCQs

ElecTech, a large electronics retailer in the Gulf region, runs its e-commerce website using a distributed system to ensure fast, reliable service for millions of customers.

- Primary Data Center (Riyadh): This is the heart of the system. It stores the main SQL database containing product inventory, user accounts, and handles all payment transactions. Any update such as a customer placing an order happens here..
- Secondary Data Center (Jeddah): This center stores a read-only version of the product catalog and caching user sessions. It also acts as a backup (failover) for the web servers if the Riyadh center becomes overloaded or experiences downtime.
- Architecture: A hybrid model
  - A centralized SQL database in Riyadh for accuracy and consistency
  - Distributed web caches for fast access to frequently viewed pages
  - A peer-to-peer Content Delivery Network(CDN) that delivers product images quickly to users
- Programming Model: The system uses Java Threads for parallel order validation.

9	The use of a peer-to-peer (P2P) CDN for distributing product images is primarily for which distributed systems goal? A) To improve scalability and reduce latency by sourcing images from geographically close peers. B) To ensure strong consistency for image updates. C) To centralize the management of all digital assets. D) To make the system more secure against external attacks.
10	The system uses Java's Thread class to validate multiple items in an order in parallel. However, when two threads simultaneously check and update the stock level for the last laptop, a race condition occurs, leading to an oversell. This problem arises because the threads are using a: A) Message Passing model, which is not suitable for this task. B) Shared Memory model without proper synchronization. C) Client-Server model, which is inherently insecure. D) Layered Architecture, which introduces too much latency.
11	The architects propose a switch from a client-server database to a structured P2P network using a DHT for inventory data to eliminate the Riyadh bottleneck. What is the most significant operational cost of this change? A) The overhead of maintaining the structured overlay and redistributing data when nodes join or leave the network. B) The DHT cannot store numerical data like inventory counts. C) The system would lose the ability to perform SQL queries. D) All inventory data would become publicly accessible.
12	The migration from a single, powerful database server to a distributed system with multiple application servers is a shift from what kind of scaling to another? A) From Horizontal Scaling to Vertical Scaling B) From Functional Scaling to Diagonal Scaling C) From Scaling Up to Scaling Down D) From Vertical Scaling to Horizontal Scaling

Question 2 [9 points - CLO 2]

Part A [3]- Select the most appropriate choice from the following MCQs

13	Why is a balanced tree generally a better data structure than a linked list for parallel operations? A) Balanced trees use less memory. B) The height of a balanced tree is O(log n), leading to fewer sequential steps in the reduction. C) Linked lists cannot be traversed in parallel. D) Balanced trees are easier to implement.
14	The "span" of a parallel algorithm, T_∞, is defined as: A) The total amount of work done by the algorithm. B) The running time using only one processor. C) The running time using an infinite number of processors. D) The ratio of work to span.
15	The Fork/Join framework in Java is particularly well-suited for which type of algorithms? A) I/O-bound tasks B) Recursive, divide-and-conquer tasks C) Simple, sequential loops D) Graphical user interface (GUI) updates
16	What is the purpose of a "sequential cutoff" in a parallel divide-and-conquer algorithm? A) To improve the accuracy of the result. B) To avoid the overhead of creating threads for small subproblems. C) To ensure tasks are processed in the correct order. D) To make the algorithm more secure.
17	Evaluating the Publish-Subscribe pattern for a system with mobile clients that frequently disconnect reveals a significant advantage. What is it? A) It provides stronger delivery guarantees than RPC. B) It eliminates the need for a network connection. C) It automatically encrypts all messages in transit. D) Its time decoupling allows publishers to send messages even when subscribers are offline.
18	In distributed systems, what is the main advantage of using indirect communication over direct RPC? A) Space and time decoupling between senders and receivers B) Better performance for small messages C) Stronger consistency guarantees D) Simpler programming model

Part B [6]- Read the scenario below and select the most appropriate answer for each of the following MCQs.

YouStream is creating a platform where users can upload their own videos. After a video is uploaded, the system must:

- Convert the video into a standard 1080p format
- Create a thumbnail image

Each of these jobs can be done separately and at the same time, so the system processes them in parallel using different services.

How the processing pipeline works:

- Upload Service: When a user uploads a video, this service receives the file and saves it in a distributed storage system.
- Dispatcher Service: This service looks at the uploaded file and creates several small tasks, i) generate the 1080p video and ii) extract the thumbnail.
- Worker Services: A group of simple, stateless workers pick up these tasks from a queue and process them. Each worker handles one task at a time, allowing many tasks to run in parallel.

19	Sarah, a performance engineer at YouStream, is optimizing the core video compression algorithm used by the Workers. The original sequential version of the algorithm takes 90 seconds to process a video. After analysis, she determines that 85% of the algorithm can be parallelized. She tests her parallel version on a machine with 16 cores. Assuming ideal conditions with perfect use of all 16 cores, what is the expected speedup using Amdahl's Law? A) 4.9x B) 5.8x C) 6.7x D) None

.

20	In reality, due to communication overhead between threads, the parallel portion only achieves an efficiency equivalent to using 12 of the 16 cores effectively. Recalculate the speedup in this case. A) 4.5x B) 5.1x C) 5.8x D) None

.

21	To handle future load, the systems architect is designing a new server fleet with 64-core machines. The goal for a key application on these servers is to achieve a speedup of at least 50 times faster than the sequential version. Using Amdahl's Law, what is the maximum percentage of the program that can remain sequential to still achieve this 50x speedup goal? A) 0.4% B) 1.3% C) 2.0% D) 3.9%

.

22	The platform must be resilient to failures. The engineering team is evaluating the semantics of their task-processing system to ensure correct behavior. For YouStream, which task is most likely to be idempotent? A) Charging a user's credit card. B) Incrementing a "total videos processed" counter. C) Overwriting a existing 1080p video file with a new encode. D) Appending a log entry with a timestamp.
23	In the communication between the Dispatcher and Workers, if a server crashes after executing a request but before replying, the client might get "at-most-once" semantics. In YouStream, this could mean: A) A video is encoded multiple times. B) A video encoding task is lost but not duplicated. C) The user is charged twice. D) The video quality is reduced.
24	A trade-off between using a single, complex multi-threaded server versus a cluster of simple, single-threaded servers, a key advantage of the cluster is: A) The overall system becomes less expensive. B) The total hardware cost is lower. C) The CPUs operate with improved cache efficiency. D) The system offers greater fault tolerance.

Hint: Amdahl's law

$$\text{Speedup} = \frac{1}{(1-P) + \frac{P}{N}}$$

Question 3 [12 points - CLO 3]
Part A [3]- Select the most appropriate choice from the following MCQs

25	In a system using Lamport Clocks for synchronization, if event A has a timestamp of 5 and event B has a timestamp of 3, what can we conclude? A. Event A definitely happened before event B. B. Event B definitely happened before event A. C. Events A and B are concurrent. D. We cannot definitively determine the order of A and B from these timestamps alone.
26	According to the FLP Impossibility result, consensus is impossible in asynchronous systems with one faulty process. How do practical systems using the Raft consensus algorithm achieve consensus? A. By assuming partial synchrony and using timeouts. B. By sacrificing safety guarantees. C. By using perfectly reliable networks. D. By avoiding asynchronous models.
27	What is the fundamental reason why the Token-Ring algorithm can suffer from high latency in large systems, even when there is low contention for the critical section? A. The token must be passed sequentially through every node in the ring, even those not interested in the resource. B. Every node must vote on whether to grant the token to a requesting node. C. A node must receive explicit permission from all other nodes before using the token. D. The token carries a large payload of system state information.
28	A time server using the Berkeley algorithm has its own clock value of 1000 at a time instance. It collects clock values from three clients: Client A: 1008 Client B: 994 Client C: 1002 What clock value will the server instruct Client B to set? A) 1001 B) 999 C) 1006 D) 1004
29	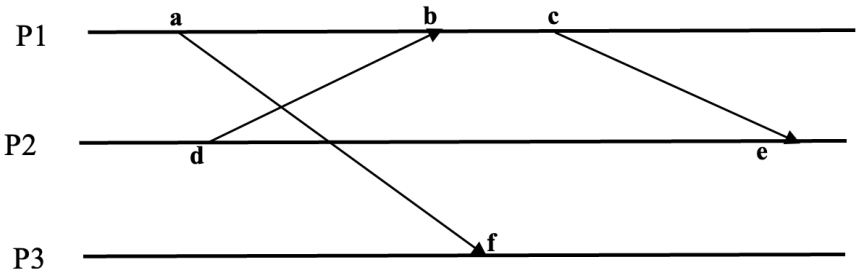 In the above figure, the vector timestamp is used. What would be the vector timestamp of event **e**. A) (3,2,1) B) (2,2,2) C) (3,2,0) D) None
30	A social media user reads a new post, then refreshes and sees an older feed. This violates: A) Linearizability B) Causal Consistency C) Monotonic Reads D) Read-your-Writes

Part B [4] - Read the following scenario and select the most appropriate choice from the following MCQs

MyTube runs a large, geographically distributed Content Delivery platform with Raft-based order replication, vector-clock event ordering, mutual-exclusion protocols for inventory, and distributed commit protocols for payments. For each of the following, provide the most appropriate answers:

31	The Order Service maintains a Raft consensus cluster (nodes A–E) to replicate customer-requests logs. During the high-traffic, the leader Node A crashes unexpectedly. After the crash, the remaining nodes each contain different portions of the replicated order log: A (crashed leader): last log **term** = 3, index = 5 B: last log term = 2, index = 4 C: last log term = 3, index = 3 D: last log term = 3, index = 4 E: last log term = 1, index = 2 Which node becomes leader for term 4? A) Node B B) Node C C) Node D D) Node E
32	When a new leader is elected in Raft, how does it ensure that the logs on all followers eventually become consistent with its own log, even if some followers have missing or extraneous entries? A. It sends a snapshot of its entire log to all followers. B. It uses the induction principle. C. It discards the logs of all followers and starts replicating from its own last committed entry. D. Followers contact other followers to avoid bottleneck at leader.
33	Meanwhile, the Notification Service uses Vector Clocks to track causal ordering of events across three internal components: • Feed Generator • Post Indexer • Delivery Queue Two important events occur: V(a) = [2,1,3] and V(b) = [1,1,4], what is the relationship between events a and b? A. a happened before b. B. b happened before a. C. a and b are concurrent. D. The clocks are inconsistent; no relationship can be determined.
34	The Inventory Service uses Ricart & Agrawala's mutual-exclusion algorithm to coordinate updates when multiple warehouses attempt to decrement stock for the same product. When a process receives a request from another process, it immediately sends a REPLY back unless: A. It is currently in the critical section or it is also requesting and its own request has a higher timestamp. B. It is not currently interested in the critical section. C. It is currently in the critical section or it is also requesting and its own request has a lower timestamp. D. It has not received a REPLY for its own request.

Part C [5]- Read the following scenario and select the most appropriate choice from the following MCQs

ShopKSA is a large e-commerce platform operating across Saudi Arabia with data centers (DC) in Riyadh (Primary) and Dammam (Secondary), connected via high-latency WAN links. The system handles inventory management, order processing, and user sessions. During the Holidays-sale (Dec 13–25, 2025), traffic surges 15 times over. At 01:39 PM on Dec 13, a fiber cut between Riyadh and Dammam causes a network partition lasting 7 minutes. Both data centers remain internally operational and continue serving local users.

- Inventory DB: Replicated using multi-master asynchronous replication (eventual consistency).
- Order Service: Writes to local Data Center, propagates via message queue.
- Session Service: Uses sessions + local cache.
- Payment Gateway: Synchronous call to external bank (High availability).

The system must decide trade-offs under Consistency-Availability-Partition (CAP) during the partition. The following questions relate to this scenario.

35	During the 3-minute partition, a user in Dammam adds the last unit of a popular smartphone to cart. Simultaneously, a Riyadh user does the same. Both see "1 left". Both checkout. Under the current scenario, which CAP guarantee is violated? A) Consistency; linearizability of inventory read/write is affected. B) Consistency; strong consistency is violated, but eventual consistency holds. C) Availability; both checkout requests must be rejected. D) Partition tolerance; system stops processing orders in one Data Center
36	The Session Service uses sticky sessions (i.e. user is tied to one server). A user in Dammam starts checkout, gets redirected to Riyadh leader (which is unavailable). System falls back to local session from a cookie and completes the checkout. This prioritizes which of the following? A) Consistency by blocking until Riyadh Data Center recovers B) Partition tolerance by merging session logs post-recovery C) Availability over Consistency of session state across Data Centers D) CAP impossibility; session state cannot be both A and C
37	The Payment Gateway requires synchronous confirmation. During partition, Dammam DC cannot reach the bank (headquarters in Riyadh). To maintain Consistency, the system queues payments and responds "Payment Pending". What do you say about this design decision? A) Chooses CP, rejects requests, so violates A. B) Maintains CA by failing over to a local payment proxy. C) Chooses AP, violates C for payment atomicity (All or none). D) Achieves CAP: Partition availability with consistency.
38	The Session Service uses local caching for user sessions. During the partition, what is the main effect on users? A) Users in Dammam can continue their sessions without interruption. B) Users are logged out immediately because session data cannot reach Riyadh. C) Session updates are lost permanently. D) Users experience inconsistent payment confirmations.
39	The Order Service writes locally and propagates via a message queue. During the partition, a user in Dammam places an order. Which statement is true? A) The order is immediately reflected in Riyadh inventory. B) The order is queued locally and will be synchronized after the partition. C) The order fails and the user is rejected. D) The order is processed twice to ensure availability.

Question 4 [5 points - CLO 4] Read the following scenario and select the most appropriate of the MCQs.

Saleh is working at SecureBank, which uses a Google File System (GFS)-style distributed file system for storing sensitive financial records and analytics data across 20 nodes in 3 data centers. Each file is divided into 64MB chunks, replicated 3 times across chunkservers. Access is controlled via Kerberos authentication, all data is AES-256 encrypted, and clients interact with a single Master Node for metadata operations. The system must tolerate node failures, network partitions, and malicious behavior while maintaining confidentiality, integrity, and availability.

40	When a client writes to a chunk in GFS, the master grants a lease to one primary chunkserver. Why is this lease mechanism critical? A. It guarantees eventual consistency across replicas. B. It ensures that all writes are serialized through the primary chunkserver. C. It allows clients to bypass the master for reads. D. It eliminates the need for replication.
41	GFS supports record append as an atomic operation. Think a GoogleDoc being edited by multiple users at the same time. Which is TRUE about this operation when multiple concurrent clients write to a record? A. All clients see data in the same order as written. B. Conflicting appends may result in partial records that clients must handle. C. The master serializes appends across all replicas using a two-phase commit. D. Record appends cannot be retried if a chunkserver crashes.
42	Which type of operation is most likely to stress the GFS master under high load? A. Large sequential reads B. File creation and metadata updates C. Direct chunk reads by clients D. Chunkserver heartbeats
43	After a client completes a write to a chunk, subsequent reads from any replica guarantee: A. Strong consistency: all reads see the latest write immediately. B. Eventual consistency: some replicas may lag, so reads can be stale briefly. C. Sequential consistency across all clients. D. Linearizability for record appends only.
44	A compromised chunkserver attempts to serve outdated chunks. How does GFS mitigate this risk? A. Primary timestamps updates; clients follow master direction for authoritative replicas B. Replication automatically ignores corrupted chunks C. Checksums are unnecessary for trusted nodes D. Clients detect malicious nodes by comparing data hashes across all replicas

Question 5 [6 points - CLO 5] Read the following scenario and select the most appropriate of the MCQs.

Fatima interns for the data analysis department of a popular local bank. Her job is to understand the customer spending behavior and look for anomalies to detect unusual activity. This requires her to download a large amount of data (>100GB) such as TransactionID, CustomerID, Amount, TransactionType, and Timestamp; from central servers to a local repository. The data is conveniently organized as key-value pairs. In the CS435 course she took at Prince Sultan University, she learned to use Hadoop for parallel processing large amounts of data. To this end, she implements two testbeds to analyze the workload:

- Centralized System: A single high-performance server (32 GB RAM, 8-core CPU) that handles all requests sequentially. The cost for this server is estimated at 12,000 SR.
- Distributed System: A cluster of 4 Raspberry Pi nodes (each with 4 GB RAM, 4-core CPU). She connects the 4 nodes and installs Hadoop on them. The cost of this setup is 2,800 SR.

To ensure reliability, she configures the system to maintain 3 replicas of each data item. She plans to run a series of tests on both setups and record the results. The following table summarizes her findings for a specific test involving mixed read/write operations.

Metric	**Centralized System (Single Node)**	**Distributed System (4 Nodes)**
Total Execution Time (sec)	120	55
Throughput (Operations/sec)	830	1818
Average Latency - Reads (ms)	5	15
Average Latency - Writes (ms)	5	90
CPU Utilization (%)	95%	45% per node
Network Messages Exchanged	N/A	~50,000
Failed Operations (due to conflicts)	0	12
Data Consistency Violations Detected	0	8
Time spent in Synchronization (%)	N/A	~25%
Leader Election Time (injected fault test)	N/A	~3.5 seconds

.

45	The distributed system shows higher throughput but worse latency. What is the BEST explanation? A. Parallelism increases per-node response times B. Network-level parallelism slows down computation C. More nodes guarantee slower writes D. Hadoop disables local caching in clusters
46	Reads are slower in the distributed system (15 ms vs 5 ms). What explains this? A. Reads require communication with multiple replicas B. Hadoop reads data only from backup nodes C. Centralized memory is slower D. Reads must pass through the leader in the centralized system

47	The throughput increases from 830 ops/sec (centralized) to 1818 ops/sec (distributed). Which principle is reflected here? A. Strong consistency increases system speed B. Adding more nodes increases aggregate processing capacity C. Removing replicas reduces overhead D. A single node is always faster than a cluster
48	What would likely happen if Fatima increases the number of nodes from 4 to 10? A. Leader election is no longer needed B. Consistency violations must drop to zero C. Synchronization overhead may increase D. All latencies must reduce
49	Fatima records that leader election takes 3.5 seconds when a fault is injected. Which component of distributed systems does this time MOST directly affect? A. Ability to maintain parallel processing B. Size of RAM needed to store data C. Performance of sequential programs D. Ability to agree on which node coordinates operations
50	One of the United Nations Sustainable Development Goals (SDGs) focuses on reducing the impacts of climate change. In the context of large-scale data processing, which statement best explains why a Hadoop cluster can be more energy-efficient per terabyte processed compared to a standalone system? A. Clusters enter low-power mode when not used. B. Parallel CPUs reduce processing time, decreasing total energy consumed for large workloads. C. Standalone systems consume more idle energy. D. Hadoop uses GPU acceleration by default for faster data processing.

-END OF EXAM-