

CHAPTER 1: INTRODUCTION TO DISTRIBUTED SYSTEM

1. Which of the following best defines a distributed system?

- A) A system where all components are located on a single machine
- B) A system in which components are located on different machines and communicate over a network
- C) A system designed to be used only on mobile devices
- D) A system with only one type of network communication protocol

Answer: B) A system in which components are located on different machines and communicate over a network

2. Which of the following is a key characteristic of a distributed system?

- A) All components are controlled by a single server
- B) Multiple processes run concurrently across different nodes
- C) There is no communication between nodes
- D) The system is monolithic and centralized

Answer: B) Multiple processes run concurrently across different nodes

3. What is the primary goal of scalability in a distributed system?

- A) To improve the speed of communication
- B) To allow the system to grow in size and handle increasing workloads
- C) To make the system more secure
- D) To ensure that the system can be accessed globally without restrictions

Answer: B) To allow the system to grow in size and handle increasing workloads

4. Which of the following is not a characteristic of a distributed system?

- A) Transparency
- B) Fault tolerance
- C) Centralized control
- D) Concurrency

Answer: C) Centralized control

5. Which design goal of a distributed system aims to ensure that the system continues to function despite failures in certain components?

- A) Efficiency
- B) Fault tolerance
- C) Scalability
- D) Heterogeneity

Answer: B) Fault tolerance

6. What does "transparency" in a distributed system mean?

- A) The system's ability to function efficiently with minimal resources
- B) The ability to protect data from unauthorized access
- C) The user is unaware of the details of the system's internal operations and component locations
- D) The ability to provide real-time data processing

Answer: C) The user is unaware of the details of the system's internal operations and component locations

7. Which of the following is a design goal of distributed systems that focuses on hiding system complexity from the user?

- A) Fault tolerance
- B) Transparency

- C) Scalability
- D) Performance

Answer: B) Transparency

8. Which type of distributed system allows resources and services to be shared across multiple machines that are physically located in different places?

- A) Client-server architecture
- B) Peer-to-peer (P2P) system
- C) Monolithic system
- D) Centralized system

Answer: B) Peer-to-peer (P2P) system

9. Which of the following is not a common type of distributed system?

- A) Client-server systems
- B) Peer-to-peer systems
- C) Hybrid systems
- D) Mainframe systems

Answer: D) Mainframe systems

10. What is the primary design goal of fault tolerance in a distributed system?

- A) To optimize network throughput
- B) To ensure continuous service even when some components fail
- C) To reduce system complexity
- D) To guarantee high-speed data processing

Answer: B) To ensure continuous service even when some components fail

CHAPTER 2: DISTRIBUTED SYSTEM ARCHITECTURE

1. Which of the following is the main purpose of a distributed system architecture?

- A) To organize the software components on a single machine
- B) To enable components to communicate and work together in a distributed environment
- C) To ensure that all components share the same memory space
- D) To control access to network resources

Answer: B) To enable components to communicate and work together in a distributed environment

2. Which of the following architecture styles in distributed systems uses a single central server to manage communication with multiple clients?

- A) Client-server architecture
- B) Peer-to-peer architecture
- C) Broker architecture
- D) Layered architecture

Answer: A) Client-server architecture

3. What is the primary role of middleware in a distributed system?

- A) To improve the efficiency of the CPU
- B) To provide communication and coordination services between distributed system components
- C) To ensure security by controlling network traffic
- D) To store data locally for faster access

Answer: B) To provide communication and coordination services between distributed system components

4. Which architecture style in distributed systems involves multiple layers of components, each layer serving a specific role or functionality?

- A) Client-server architecture
- B) Peer-to-peer architecture
- C) Layered architecture
- D) Hybrid architecture

Answer: C) Layered architecture

5. In a client-server architecture, which component typically acts as the "server"?

- A) The entity that makes requests for services
- B) The entity that provides resources or services to clients
- C) The component that manages data consistency
- D) The entity that stores data locally

Answer: B) The entity that provides resources or services to clients

6. Which middleware service allows different components of a distributed system to communicate by using messages or events?

- A) Database middleware
- B) Message-oriented middleware (MOM)
- C) Object request broker (ORB)
- D) Remote procedure call (RPC) middleware

Answer: B) Message-oriented middleware (MOM)

7. Which of the following is not a characteristic of middleware in a distributed system?

- A) Enables communication between different system components
- B) Provides services such as authentication and data encryption
- C) Acts as an interface between different hardware components
- D) Manages system-wide resources like CPU and memory

Answer: D) Manages system-wide resources like CPU and memory

8. Which distributed system architecture style involves a decentralized approach where each node can act as both a client and a server?

- A) Client-server architecture
- B) Peer-to-peer architecture
- C) Event-based architecture
- D) Microservices architecture

Answer: B) Peer-to-peer architecture

9. In the broker architecture style, what is the primary role of the broker?

- A) To process requests from clients and direct them to the appropriate service provider
- B) To manage the database of all connected clients
- C) To establish a secure connection between servers
- D) To cache results from previous computations to improve performance

Answer: A) To process requests from clients and direct them to the appropriate service provider

10. Which type of distributed system architecture allows different services to be developed and deployed independently but communicates through well-defined interfaces?

- A) Microservices architecture
- B) Client-server architecture
- C) Peer-to-peer architecture
- D) Layered architecture

Answer: A) Microservices architecture

11. Which of the following is an example of object-based middleware that allows the communication between distributed objects?

- A) Remote Procedure Call (RPC)
- B) Object Request Broker (ORB)
- C) Message Queueing (MQ)
- D) File Transfer Protocol (FTP)

Answer: B) Object Request Broker (ORB)

12. Which distributed system architecture style uses separate and independent layers to implement specific services such as data management, presentation, and business logic?

- A) Layered architecture
- B) Event-driven architecture
- C) Microservices architecture
- D) Client-server architecture

Answer: A) Layered architecture

13. Which of the following is a key advantage of middleware in distributed systems?

- A) It ensures that all components use the same operating system
- B) It abstracts the underlying hardware and software complexity, allowing

seamless communication

- C) It stores data centrally for quick access
- D) It directly manages network bandwidth

Answer: B) It abstracts the underlying hardware and software complexity, allowing seamless communication

14. Which architecture style focuses on the independent deployment and operation of services, which communicate via lightweight protocols, often used for scalability and flexibility?

- A) Layered architecture
- B) Microservices architecture
- C) Client-server architecture
- D) Peer-to-peer architecture

Answer: B) Microservices architecture

15. Which of the following is true about peer-to-peer (P2P) architecture?

- A) There is a central server that manages all communication between nodes
- B) Each node in the system can act as both a client and a server
- C) It relies heavily on a layered architecture to function
- D) It requires middleware for basic communication between components

Answer: B) Each node in the system can act as both a client and a server

CHAPTER 3: PROCESS AND THREADS

1. What is the primary difference between a process and a thread in a distributed system?

- A) A thread has its own memory space, whereas a process shares memory with other threads
- B) A process is a unit of execution with its own memory space, whereas a thread is

a smaller unit of execution within a process

C) A process can only execute one task at a time, whereas a thread can execute multiple tasks

D) A thread operates in user mode, whereas a process operates in kernel mode

Answer: B) A process is a unit of execution with its own memory space, whereas a thread is a smaller unit of execution within a process

2. Which of the following is an example of concurrent execution in a distributed system?

A) One process waiting for the other to complete before starting

B) Multiple threads within a single process executing independently

C) A single-threaded process that executes one instruction at a time

D) A process executing without utilizing the CPU

Answer: B) Multiple threads within a single process executing independently

3. In virtualization, what is the role of a hypervisor?

A) It manages memory allocation for processes within a virtual machine

B) It provides a platform for multiple virtual machines to share physical resources

C) It stores data in a virtualized database

D) It ensures that a system's hardware is shared equally among processes

Answer: B) It provides a platform for multiple virtual machines to share physical resources

4. Which of the following is not a benefit of virtualization in distributed systems?

A) Improved resource utilization

B) Isolation of applications from the underlying hardware

- C) Ability to run multiple operating systems simultaneously on the same physical machine
- D) Increased system complexity and overhead

Answer: D) Increased system complexity and overhead

5. In a client-server architecture, which of the following is true?

- A) Both the client and the server can initiate communication independently
- B) The client requests services from the server, which then responds to the client's request
- C) The client and the server cannot communicate unless both are on the same machine
- D) Clients store all data and manage resources while the server simply communicates with them

Answer: B) The client requests services from the server, which then responds to the client's request

6. What does code migration refer to in a distributed system?

- A) The process of moving an entire database from one server to another
- B) The ability to move code execution from one machine to another during runtime
- C) The process of transferring code from one software library to another
- D) The relocation of data within the system

Answer: B) The ability to move code execution from one machine to another during runtime

7. Which of the following is a key advantage of using threads within a process?

- A) Threads cannot share memory, which makes them more isolated
- B) Threads provide independent execution but share the same memory space, making communication faster
- C) Threads are always executed sequentially and in isolation
- D) Threads consume more resources than processes due to their memory isolation

Answer: B) Threads provide independent execution but share the same memory space, making communication faster

8. What is the main goal of virtualization in distributed systems?

- A) To provide a real-time operating system
- B) To allow multiple applications to run on the same physical machine with different operating systems
- C) To make distributed systems more fault-tolerant by using redundant hardware
- D) To make it easier to manage network protocols between servers

Answer: B) To allow multiple applications to run on the same physical machine with different operating systems

9. In client-server architecture, what role does the server play?

- A) The server requests resources from the client
- B) The server stores data and provides resources or services to clients
- C) The server processes the data and forwards it to clients for display
- D) The server only handles database queries and nothing else

Answer: B) The server stores data and provides resources or services to clients

10. Which of the following is an advantage of code migration in distributed systems?

- A) It reduces the need for virtual machines
- B) It allows computation to move closer to the data, potentially improving performance
- C) It makes it impossible to transfer large datasets
- D) It simplifies the system by ensuring all tasks are processed on a single server

Answer: B) It allows computation to move closer to the data, potentially improving performance

11. Which of the following best defines the role of virtual machines (VMs) in virtualization?

- A) They are the physical hardware components that host the operating systems
- B) They are software environments that simulate hardware to allow multiple operating systems to run on the same machine
- C) They only manage memory for processes in a system
- D) They allow multiple processes within a single operating system to share the same memory space

Answer: B) They are software environments that simulate hardware to allow multiple operating systems to run on the same machine

12. In the context of code migration, what is a major challenge?

- A) Ensuring that the migrated code is compatible with the new environment
- B) Making sure that the code is always executed on the same machine
- C) Moving large amounts of data between different parts of the system
- D) Ensuring that every client in the system receives the same code

Answer: A) Ensuring that the migrated code is compatible with the new environment

13. Which of the following best describes the concept of threads in distributed systems?

- A) A thread is a collection of processes that share the same memory space
- B) A thread is an individual unit of execution within a process that shares the process's memory space
- C) Threads cannot communicate with each other because they are isolated
- D) Threads are separate programs that do not interact with each other

Answer: B) A thread is an individual unit of execution within a process that shares the process's memory space

14. In a client-server model, which of the following is typically handled by the client?

- A) Managing the server's storage resources
- B) Making requests for services and displaying results to the user
- C) Maintaining a persistent connection with the database
- D) Providing resources to multiple clients

Answer: B) Making requests for services and displaying results to the user

15. Which virtualization technique allows multiple virtual machines to run on a single physical machine?

- A) Full virtualization
- B) Partial virtualization
- C) Containerization
- D) Client-server architecture

Answer: A) Full virtualization

CHAPTER 4: COMMUNICATION

1. Which of the following best describes communication in distributed systems?

- A) The process of transferring data between a client and a server in a single machine
- B) The exchange of data between processes or nodes in a distributed system
- C) The execution of code on a single machine
- D) The storage of data in a centralized database

Answer: B) The exchange of data between processes or nodes in a distributed system

2. What is the primary purpose of RPC (Remote Procedure Call) in distributed systems?

- A) To allow a program to make requests for services on a remote server as if they were local function calls
- B) To enable message-based communication between nodes
- C) To handle data storage in a distributed environment
- D) To perform real-time data synchronization between nodes

Answer: A) To allow a program to make requests for services on a remote server as if they were local function calls

3. In Message-Oriented Communication, what is the typical medium used to send messages between distributed components?

- A) Shared memory
- B) Direct function calls
- C) Queues or message brokers
- D) TCP/IP connections

Answer: C) Queues or message brokers

4. Which of the following is an advantage of Message-Oriented Communication in distributed systems?

- A) Faster response times
- B) Simpler application logic
- C) Loose coupling and asynchronous communication between components
- D) Easier to maintain due to centralized data storage

Answer: C) Loose coupling and asynchronous communication between components

5. What is multicasting communication in a distributed system?

- A) Sending a message to all nodes in the system
- B) Sending a message to a single node
- C) Sending a message to multiple specific recipients simultaneously
- D) Sending a message over a unidirectional communication channel

Answer: C) Sending a message to multiple specific recipients simultaneously

6. Which of the following protocols is typically used for multicasting communication in a distributed system?

- A) HTTP
- B) FTP
- C) IP Multicast
- D) SMTP

Answer: C) IP Multicast

7. Which of the following is a disadvantage of using RPC (Remote Procedure Call) in distributed systems?

- A) It simplifies the communication between distributed components by abstracting remote interactions
- B) It introduces latency due to the network round-trip
- C) It is highly secure and prevents data leakage
- D) It allows easy handling of network failures

Answer: B) It introduces latency due to the network round-trip

8. What is the main advantage of using Message-Oriented Middleware (MOM) in distributed systems?

- A) It allows components to send messages synchronously
- B) It enables asynchronous communication and decouples message producers from consumers
- C) It ensures that all messages are delivered in real-time
- D) It increases the system's performance by reducing the need for message acknowledgment

Answer: B) It enables asynchronous communication and decouples message producers from consumers

9. In multicast communication, which of the following is a challenge?

- A) Ensuring that all messages are received by every node
- B) Ensuring that a message is only delivered to one recipient
- C) Maintaining the order of messages across nodes
- D) Handling message acknowledgment between multiple recipients

Answer: A) Ensuring that all messages are received by every node

10. Which of the following is a typical use case for RPC (Remote Procedure Call) in distributed systems?

- A) Sharing resources between peers in a P2P network
- B) Enabling remote method invocation on a server from a client
- C) Sending broadcast messages to all clients in a system
- D) Transferring files between distributed nodes

Answer: B) Enabling remote method invocation on a server from a client

11. How does Message-Oriented Communication handle message delivery?

- A) It guarantees the delivery of messages in real-time
- B) It uses a pull-based model to retrieve messages at the client's request
- C) It relies on a store-and-forward approach to deliver messages asynchronously
- D) It requires both sender and receiver to be connected at the same time

Answer: C) It relies on a store-and-forward approach to deliver messages asynchronously

12. Which of the following is true about RPC in distributed systems?

- A) RPC always sends messages in a non-blocking manner
- B) RPC requires both the client and server to be written in the same programming language
- C) RPC abstracts the network communication details, allowing developers to call remote services as if they were local
- D) RPC is only used for file sharing across distributed nodes

Answer: C) RPC abstracts the network communication details, allowing developers to call remote services as if they were local

13. What distinguishes multicasting communication from broadcasting in distributed systems?

- A) Multicasting sends messages to a specific group of receivers, while broadcasting sends messages to all nodes in the system
- B) Multicasting requires a connection between all nodes, while broadcasting does not
- C) Multicasting is used for asynchronous communication, while broadcasting is synchronous
- D) Multicasting is a point-to-point communication, while broadcasting is a one-to-many communication

Answer: A) Multicasting sends messages to a specific group of receivers, while broadcasting sends messages to all nodes in the system

14. Which of the following is not typically used in Message-Oriented Communication?

- A) Message queues
- B) Publish-subscribe models
- C) Remote procedure calls
- D) Asynchronous message delivery

Answer: C) Remote procedure calls

15. What is one of the key benefits of multicasting in a distributed system?

- A) It reduces bandwidth usage when sending data to multiple receivers
- B) It guarantees message delivery to every node in the network
- C) It only supports one-to-one communication
- D) It simplifies the process of managing remote procedure calls

Answer: A) It reduces bandwidth usage when sending data to multiple receivers

CHAPTER 5: NAMING

1. Which of the following refers to a unique identifier assigned to each entity in a distributed system?

- A) Address
- B) Name
- C) Identifier
- D) Attribute

Answer: C) Identifier

2. In a distributed system, which of the following is used to refer to the location of an entity or resource in the network?

- A) Name
- B) Identifier
- C) Address
- D) Attribute

Answer: C) Address

3. Structured naming in a distributed system is useful for:

- A) Organizing data hierarchically
- B) Uniquely identifying an entity
- C) Assigning specific attributes to entities
- D) Providing a single address for the system

Answer: A) Organizing data hierarchically

4. Which of the following is true about attribute naming in distributed systems?

- A) Attribute names are unique only within a given context or object.
- B) Attribute names provide a way to uniquely identify an entity across the entire system.

- C) Attribute names are the same as identifiers in a distributed system.
- D) Attribute names are used to assign network addresses.

Answer: A) Attribute names are unique only within a given context or object.

5. A distributed system uses structured naming for which of the following reasons?

- A) To simplify the address resolution process
- B) To associate a unique address with a specific entity
- C) To represent entities in a logical hierarchy or namespace
- D) To assign unique identifiers to resources

Answer: C) To represent entities in a logical hierarchy or namespace

6. What is the primary role of an identifier in a distributed system?

- A) To locate an entity in the system's network
- B) To represent a set of attributes associated with an entity
- C) To uniquely recognize an entity within the system
- D) To define the network protocol for communication

Answer: C) To uniquely recognize an entity within the system

7. In a distributed system, what does the address refer to?

- A) A physical or logical location of an entity
- B) A label for an entity
- C) A reference to an entity's attributes
- D) A protocol used for communication

Answer: A) A physical or logical location of an entity

8. In the context of distributed systems, which of the following terms refers to a more complex naming scheme that includes namespaces and hierarchies?

- A) Flat naming
- B) Structured naming
- C) Attribute naming
- D) Identifiers

Answer: B) Structured naming

9. Which of the following can be an example of an attribute in a distributed system?

- A) A node's physical memory size
- B) A unique identifier assigned to a node
- C) The IP address of a node
- D) A resource's network address

Answer: A) A node's physical memory size

10. What is the primary benefit of structured naming in a distributed system?

- A) It allows for efficient querying of resources.
- B) It simplifies addressing and makes the network scalable.
- C) It ensures that all nodes have the same identifier.
- D) It guarantees that no conflicts will occur in the system.

Answer: B) It simplifies addressing and makes the network scalable.

CHAPTER 6 : COORDINATION

1. What is the main purpose of clock synchronization in distributed systems?

- A) To ensure all distributed nodes have a unique identifier
- B) To make sure all nodes in a distributed system share a consistent view of time

- C) To prevent nodes from communicating with each other
- D) To synchronize data across all nodes in a system

Answer: B) To make sure all nodes in a distributed system share a consistent view of time

2. Which of the following is not a type of clock synchronization algorithm?

- A) Lamport's logical clock
- B) Berkeley algorithm
- C) NTP (Network Time Protocol)
- D) Chandy-Lamport algorithm

Answer: D) Chandy-Lamport algorithm

3. Which of the following best defines a logical clock in a distributed system?

- A) A physical clock that is synchronized between distributed nodes
- B) A clock used to simulate time in a distributed system, which may not correspond to real time
- C) A clock used only for event sequencing in a single machine
- D) A clock that only tracks system resource usage in a distributed system

Answer: B) A clock used to simulate time in a distributed system, which may not correspond to real time

4. What is the primary purpose of mutual exclusion in distributed systems?

- A) To ensure that multiple processes can access a resource simultaneously
- B) To allow multiple processes to work on different parts of the system at the same time
- C) To ensure that only one process at a time can access a critical section or shared

resource

D) To synchronize clocks across all processes in a distributed system

Answer: C) To ensure that only one process at a time can access a critical section or shared resource

5. Which of the following is the main goal of an election algorithm in distributed systems?

A) To synchronize the clocks across all nodes

B) To choose a coordinator or leader among a group of distributed processes

C) To ensure mutual exclusion among distributed nodes

D) To perform fault tolerance in case of node failures

Answer: B) To choose a coordinator or leader among a group of distributed processes

6. In the Lamport's Logical Clock algorithm, what is the basic idea behind the increment of the logical clock?

A) The logical clock is incremented every time an event occurs in the system

B) The logical clock is incremented based on the synchronization of physical clocks

C) The logical clock is incremented only when a process sends a message

D) The logical clock is never incremented

Answer: A) The logical clock is incremented every time an event occurs in the system

7. In mutual exclusion algorithms for distributed systems, which of the following is true?

- A) Mutual exclusion can be easily achieved by having a central coordinator
- B) Every process must be aware of all other processes in the system
- C) Distributed mutual exclusion algorithms generally use a centralized clock
- D) Mutual exclusion algorithms allow multiple processes to access the critical section simultaneously

Answer: B) Every process must be aware of all other processes in the system

8. In the Ring-based Election Algorithm, what is the primary action taken by processes?

- A) A process sends its ID to the next process in the ring and waits for a decision
- B) A process selects the leader from its neighbors based on a priority queue
- C) A process randomly selects a leader among all processes
- D) A process informs the central coordinator about its state

Answer: A) A process sends its ID to the next process in the ring and waits for a decision

9. Gossip-based coordination in distributed systems is used primarily for which of the following?

- A) Synchronizing clocks between all nodes in the system
- B) Broadcasting messages to a small set of recipients
- C) Disseminating information or updates in an epidemic-like manner to a subset of nodes
- D) Choosing a leader among distributed processes

Answer: C) Disseminating information or updates in an epidemic-like manner to a subset of nodes

10. What is the main benefit of gossip-based coordination in a distributed system?

- A) It requires minimal overhead for communication between nodes
- B) It guarantees the reliability of message delivery
- C) It ensures that all nodes are synchronized in real-time
- D) It is useful for high-priority, urgent message passing only

Answer: A) It requires minimal overhead for communication between nodes

11. Which of the following is an advantage of logical clocks over physical clocks in distributed systems?

- A) Logical clocks are more accurate and closer to real-world time
- B) Logical clocks do not require synchronization between different machines
- C) Logical clocks can help in maintaining the same real-time view across all nodes
- D) Logical clocks are synchronized by a global time server

Answer: B) Logical clocks do not require synchronization between different machines

12. Which of the following is a commonly used algorithm for mutual exclusion in distributed systems?

- A) Lamport's algorithm
- B) Ricart-Agrawala algorithm
- C) Chandy-Lamport algorithm
- D) Both A and B

Answer: D) Both A and B

13. What is the key idea behind Lamport's logical clock algorithm in distributed systems?

- A) Clocks are synchronized based on messages exchanged between nodes
- B) Each process maintains a local clock and increments it whenever an event

occurs

- C) Logical clocks use real-time physical clocks for synchronization
- D) Lamport's clock does not provide any guarantees about the order of events

Answer: B) Each process maintains a local clock and increments it whenever an event occurs

14. Which of the following is a key feature of gossip protocols in distributed systems?

- A) High consistency in data across all nodes
- B) Acknowledgments for each message received by every node
- C) Propagation of updates or information in a decentralized, epidemic-like manner
- D) Strict control over message delivery and prioritization

Answer: C) Propagation of updates or information in a decentralized, epidemic-like manner

15. What does the Berkeley algorithm for clock synchronization aim to achieve in a distributed system?

- A) Synchronizing all nodes to a central time server
- B) Providing global time synchronization using a master node and adjustment of local clocks
- C) Making sure that clocks in all nodes are perfectly synchronized
- D) Using logical clocks to synchronize events across processes

Answer: B) Providing global time synchronization using a master node and adjustment of local clocks

CHAPTER 7: REPLICATION AND CONSISTENCY

1. What is the primary goal of replication in distributed systems?

- A) To reduce the time required for data processing
- B) To maintain multiple copies of data across different nodes for fault tolerance and availability
- C) To increase network bandwidth
- D) To minimize the number of requests sent to the central server

Answer: B) To maintain multiple copies of data across different nodes for fault tolerance and availability

2. Which of the following best describes data-centric consistency in a distributed system?

- A) It focuses on keeping the client experience consistent, regardless of the underlying data state
- B) It ensures that updates to a particular piece of data are consistently reflected across all replicas
- C) It prioritizes client requests over server-side updates
- D) It guarantees high availability by allowing data replicas to be inconsistent

Answer: B) It ensures that updates to a particular piece of data are consistently reflected across all replicas

3. Which of the following is the focus of client-centric consistency models?

- A) Ensuring that the client always sees the most up-to-date version of the data
- B) Maintaining consistency between multiple data replicas in the system
- C) Enforcing the order of writes at the server level
- D) Ensuring that all data replicas are identical across all clients

Answer: A) Ensuring that the client always sees the most up-to-date version of the data

4. Which of the following is a primary challenge when using replication in distributed systems?

- A) Ensuring that all clients can update the data simultaneously
- B) Managing the trade-off between consistency and availability during data replication
- C) Eliminating the need for consistency protocols
- D) Making replicas globally consistent at all times

Answer: B) Managing the trade-off between consistency and availability during data replication

5. Which of the following consistency models ensures that all clients observe the same sequence of operations in the same order?

- A) Eventual consistency
- B) Strong consistency
- C) Causal consistency
- D) FIFO consistency

Answer: D) FIFO consistency

6. In replica management, what does quorum-based replication refer to?

- A) Using a single replica as the authoritative source of truth
- B) A method where a certain number of nodes must agree on a value before an update is considered successful
- C) Allowing clients to update any replica without synchronization
- D) Replicating data across all available nodes

Answer: B) A method where a certain number of nodes must agree on a value before an update is considered successful

7. Which of the following consistency protocols is commonly used to achieve strong consistency in a distributed system?

- A) Two-Phase Commit (2PC)
- B) Paxos
- C) Raft
- D) All of the above

Answer: D) All of the above

8. What is the primary goal of eventual consistency in distributed systems?

- A) To ensure that all data replicas are updated at the same time
- B) To guarantee that all clients always see the same data
- C) To allow replicas to temporarily diverge, but eventually converge to a consistent state
- D) To prioritize consistency over availability

Answer: C) To allow replicas to temporarily diverge, but eventually converge to a consistent state

9. What is a key advantage of using data-centric consistency models?

- A) They allow clients to access data faster by caching copies locally
- B) They prioritize availability over consistency, improving fault tolerance
- C) They provide better guarantees of consistency across replicas, ensuring a uniform state
- D) They make it easier to manage client-side states independently of data updates

Answer: C) They provide better guarantees of consistency across replicas, ensuring a uniform state

10. Which of the following consistency protocols requires a leader node to ensure consistency across replicas in a distributed system?

- A) Paxos
- B) Raft
- C) Vector Clocks
- D) Causal Consistency

Answer: B) Raft

11. Which of the following is an example of a client-centric consistency model?

- A) Strong consistency
- B) Eventual consistency
- C) Causal consistency
- D) Linearizability

Answer: C) Causal consistency

12. What is the main benefit of replicating data across multiple nodes in a distributed system?

- A) It makes data access slower by introducing network delays
- B) It increases system availability and fault tolerance by providing multiple data copies
- C) It ensures that all nodes in the system are synchronized at all times
- D) It reduces the need for consistency protocols

Answer: B) It increases system availability and fault tolerance by providing multiple data copies

13. Which of the following replication techniques ensures that all replicas have the same data at all times, even in the presence of network failures?

- A) Synchronous replication
- B) Asynchronous replication
- C) Lazy replication
- D) Eventual replication

Answer: A) Synchronous replication

14. In the context of consistency protocols, what does linearizability guarantee?

- A) It guarantees that updates are eventually consistent across replicas
- B) It ensures that all operations appear to happen instantaneously at some point between their start and end
- C) It ensures the most recent data is always available to clients
- D) It ensures that all nodes have a consistent view of the system's state

Answer: B) It ensures that all operations appear to happen instantaneously at some point between their start and end

15. Which of the following best describes replica management in distributed systems?

- A) The process of synchronizing all data replicas across different nodes to ensure uniformity
- B) The process of selecting a primary replica based on the node's processing power
- C) The automatic replication of data between clients and servers
- D) The creation of multiple replicas of data without concern for consistency

Answer: A) The process of synchronizing all data replicas across different nodes to ensure uniformity

CHAPTER 8: FAULT TOLERANCE

1. Which of the following describes fault tolerance in a distributed system?

- a) Ensuring data integrity through encryption
- b) Preventing system failures by avoiding crashes
- c) The ability of a system to continue operating despite failures
- d) Making a system operate faster under load

Answer: c) The ability of a system to continue operating despite failures

2. In the context of recovery in distributed systems, what does the term "rollback" refer to?

- a) Undoing changes made after a failure
- b) Moving forward to a new state
- c) Starting the system from the last checkpoint
- d) Storing system data periodically

Answer: a) Undoing changes made after a failure

3. Which of the following is a key requirement for implementing distributed commit protocols?

- a) Coordination among multiple nodes to agree on a decision
- b) Quick recovery from crashes
- c) Redundancy of all data
- d) Uniformity of system architecture

Answer: a) Coordination among multiple nodes to agree on a decision

4. What is the primary purpose of checkpoints in a distributed system?

- a) To store backup copies of data
- b) To create a consistent system state that can be restored in case of failure

- c) To distribute tasks among multiple nodes
- d) To synchronize the clock across different machines

Answer: b) To create a consistent system state that can be restored in case of failure

5. In the Two-Phase Commit (2PC) protocol, what happens if the coordinator crashes after sending the "commit" message but before the participants send their acknowledgment?

- a) The system enters an inconsistent state
- b) The participants automatically assume the transaction is aborted
- c) The system waits for the coordinator to recover and send a final decision
- d) The transaction is automatically rolled back

Answer: c) The system waits for the coordinator to recover and send a final decision

6. Which of the following describes a key challenge in implementing fault tolerance in distributed systems?

- a) Ensuring there is always a single point of failure
- b) Achieving data consistency while allowing partial system failures
- c) Limiting the number of nodes in the system
- d) Ensuring that all nodes are identical in architecture

Answer: b) Achieving data consistency while allowing partial system failures

7. In distributed systems, which of the following recovery techniques helps in minimizing the amount of data that needs to be reprocessed after a failure?

- a) Redundancy
- b) Rollback

- c) Checkpointing
- d) Replication

Answer: c) Checkpointing

8. Which of the following protocols ensures that either all participants in a distributed transaction commit or none do (Atomicity)?

- a) Paxos
- b) Raft
- c) Two-Phase Commit (2PC)
- d) Three-Phase Commit (3PC)

Answer: c) Two-Phase Commit (2PC)

9. What is the main issue with the Two-Phase Commit (2PC) protocol?

- a) Lack of fault tolerance
- b) Possibility of deadlock in case of crash failures
- c) Complexity in ensuring consistency
- d) It is not suitable for small distributed systems

Answer: b) Possibility of deadlock in case of crash failures

10. What is the purpose of "write-ahead logging" in the context of fault tolerance?

- a) To ensure that logs are written to disk before changes are made to the data
- b) To record the changes that have been made to a system for performance optimization
- c) To log only committed transactions
- d) To record the failure events for analysis

Answer: a) To ensure that logs are written to disk before changes are made to the data

CHAPTER 9: SECURITY

1. Which of the following is the primary goal of security in distributed systems?

- A) To reduce the overall network latency
- B) To ensure data confidentiality, integrity, and availability
- C) To maximize system throughput
- D) To ensure only authorized users can access the system resources

Answer: B) To ensure data confidentiality, integrity, and availability

2. Which of the following is a common threat in distributed systems?

- A) Denial of Service (DoS) attacks
- B) Memory overflow errors
- C) Increased throughput
- D) Load balancing issues

Answer: A) Denial of Service (DoS) attacks

3. What does access control in distributed systems primarily focus on?

- A) Ensuring that data is always available to the user
- B) Restricting unauthorized users or processes from accessing system resources
- C) Encrypting data while it is being transmitted
- D) Synchronizing all clocks across nodes

Answer: B) Restricting unauthorized users or processes from accessing system resources

4. In distributed systems, which of the following threats could lead to unauthorized access to resources?

- A) Spoofing
- B) Data redundancy
- C) Load balancing
- D) Synchronization

Answer: A) Spoofing

5. What is the primary purpose of encryption in secure communication channels within a distributed system?

- A) To improve network speed
- B) To ensure data confidentiality during transmission
- C) To ensure that data is available to all nodes
- D) To guarantee that the sender and receiver are synchronized

Answer: B) To ensure data confidentiality during transmission

6. Which of the following is an example of an access control mechanism used in distributed systems?

- A) Public key infrastructure (PKI)
- B) Firewalls
- C) Role-based access control (RBAC)
- D) Data encryption

Answer: C) Role-based access control (RBAC)

7. Which of the following best describes authentication in distributed systems?

- A) Verifying the identity of a user or system
- B) Ensuring that a user has the correct access privileges
- C) Encrypting communication between systems
- D) Ensuring that data is not altered during transmission

Answer: A) Verifying the identity of a user or system

8. Which type of attack aims to overload a system and make it unavailable to users in a distributed system?

- A) Man-in-the-middle attack
- B) Denial of Service (DoS) attack
- C) Phishing attack
- D) Replay attack

Answer: B) Denial of Service (DoS) attack

9. In the context of secure channels, what is the primary role of a secure channel in a distributed system?

- A) To ensure that messages are transmitted over an encrypted and authenticated path between two parties
- B) To distribute the load evenly among nodes
- C) To store messages temporarily before they are processed
- D) To ensure that the data is not cached

Answer: A) To ensure that messages are transmitted over an encrypted and authenticated path between two parties

10. Which of the following is a technique used to ensure data integrity in distributed systems?

- A) Public key encryption
- B) Digital signatures
- C) Load balancing
- D) Network address translation (NAT)

Answer: B) Digital signatures

11. Which of the following threats is mitigated by firewalls in distributed systems?

- A) Spoofing
- B) Man-in-the-middle attack
- C) Unauthorized access
- D) Data leakage

Answer: C) Unauthorized access

12. Which of the following is a characteristic of symmetric encryption in distributed systems?

- A) The same key is used for both encryption and decryption
- B) The encryption key is public, and the decryption key is private
- C) It is faster than asymmetric encryption
- D) Both A and C are correct

Answer: D) Both A and C are correct

13. In the context of access control, what does least privilege mean?

- A) Giving users only the permissions necessary to perform their tasks
- B) Granting all users full access to all system resources
- C) Allowing administrators to have unrestricted access to all data
- D) Ensuring that only trusted users are allowed to log in

Answer: A) Giving users only the permissions necessary to perform their tasks

14. Which of the following is the primary concern of confidentiality in distributed systems?

- A) Ensuring that only authorized users can access sensitive data
- B) Ensuring that all data is available to all users
- C) Ensuring that data is consistent across all nodes
- D) Ensuring that data can be quickly recovered in case of failure

Answer: A) Ensuring that only authorized users can access sensitive data

15. Which of the following is a secure communication protocol often used in distributed systems to provide encryption and authentication?

- A) HTTP
- B) SSL/TLS
- C) FTP
- D) SMTP

Answer: B) SSL/TLS

1. Which of the following best describes a confidentiality threat in a distributed system?

- A) Unauthorized modification of data by an attacker
- B) An attacker intercepting and reading sensitive data during transmission
- C) A system crash causing data loss
- D) Unauthorized access to system resources by a legitimate user

Answer: B) An attacker intercepting and reading sensitive data during transmission

2. Which of the following is an example of an integrity threat in a distributed system?

- A) A man-in-the-middle attack altering messages between two communicating systems
- B) A distributed denial of service (DDoS) attack that causes system unavailability
- C) A buffer overflow that leads to a system crash
- D) A phishing attack that tricks users into revealing passwords

Answer: A) A man-in-the-middle attack altering messages between two communicating systems

3. What is the primary goal of availability in the context of security?

- A) Ensuring that only authorized users can access the system
- B) Protecting the integrity of data from unauthorized changes
- C) Ensuring that services and resources are available when needed, even under attack
- D) Ensuring that data is encrypted during transmission

Answer: C) Ensuring that services and resources are available when needed, even under attack

4. Which of the following attacks is primarily a confidentiality threat in distributed systems?

- A) Distributed denial-of-service (DDoS)
- B) Eavesdropping or sniffing of network traffic
- C) Buffer overflow
- D) Data manipulation or tampering

Answer: B) Eavesdropping or sniffing of network traffic

5. Which of the following is a common threat to integrity in a distributed system?

- A) Data being modified by unauthorized users or systems
- B) Data being completely destroyed or lost
- C) The system crashing and preventing access to resources
- D) A firewall blocking legitimate access requests

Answer: A) Data being modified by unauthorized users or systems

6. What is the primary concern of confidentiality in security?

- A) Ensuring that the data is accurate and unchanged
- B) Preventing unauthorized access to sensitive information
- C) Ensuring that the system is available even under stress
- D) Preventing denial-of-service attacks

Answer: B) Preventing unauthorized access to sensitive information

7. Which of the following is an example of an availability threat in distributed systems?

- A) A user gaining unauthorized access to system data
- B) An attacker flooding a system with traffic, causing it to crash (DDoS)
- C) A malicious actor modifying data stored on a server
- D) An attacker intercepting and reading data during communication

Answer: B) An attacker flooding a system with traffic, causing it to crash (DDoS)

8. Which of the following attacks could be considered a confidentiality threat to a distributed system?

- A) SQL injection
- B) Session hijacking
- C) Data alteration via man-in-the-middle attack
- D) Phishing for user credentials

Answer: B) Session hijacking

9. Which of the following is not a typical availability threat?

- A) Denial of Service (DoS) attacks
- B) Server crashes due to software bugs
- C) Unauthorized modification of system data
- D) Resource exhaustion caused by excessive requests

Answer: C) Unauthorized modification of system data

10. Which of the following best defines integrity in the context of security?

- A) The system is only accessible to authorized users
- B) The system is always available when needed
- C) The data remains accurate, unaltered, and trustworthy
- D) Data is protected from unauthorized access and eavesdropping

Answer: C) The data remains accurate, unaltered, and trustworthy

11. An attacker gaining access to a database and modifying records without authorization is an example of which type of security threat?

- A) Confidentiality threat
- B) Integrity threat
- C) Availability threat
- D) Authentication threat

Answer: B) Integrity threat

12. Which of the following best describes a confidentiality threat in terms of data access?

- A) An attacker stealing encryption keys to decrypt sensitive data
- B) An attacker attempting to destroy all system resources
- C) A user misusing their access privileges to alter data
- D) An attacker flooding the network with data packets to prevent access

Answer: A) An attacker stealing encryption keys to decrypt sensitive data

13. A denial-of-service (DoS) attack affects which of the following aspects of a distributed system?

- A) Confidentiality
- B) Integrity
- C) Availability
- D) Authentication

Answer: C) Availability

14. Which of the following describes a distributed denial-of-service (DDoS) attack as an availability threat?

- A) It floods the target system with excessive requests, causing the service to become unavailable to legitimate users.
- B) It alters data being transmitted between systems.
- C) It allows unauthorized users to access sensitive data.
- D) It modifies system configurations to compromise the data's accuracy.

Answer: A) It floods the target system with excessive requests, causing the service to become unavailable to legitimate users.

15. To mitigate integrity threats in a distributed system, which of the following techniques is most commonly used?

- A) Public key encryption
- B) Digital signatures and hashing
- C) Load balancing
- D) IP filtering

Answer: B) Digital signatures and hashing