


CS 417 – DISTRIBUTED SYSTEMS

Week 5: Exam 1 Review

Paul Krzyzanowski



© 2022 Paul Krzyzanowski. No part of this content may be reproduced or reposted in whole or in part in any manner without the permission of the copyright owner.

Question 1

In distributed systems, a *single system image*:

- (a) Runs a single operating system that is distributed across multiple computers.
- (b) Comprises multiple computers, each with an identical architecture and operating system.
- (c) Runs replicated copies of an application across multiple computers for high availability.
- (d) Gives users the illusion that they are interacting with one system.

Question 2

Suppose an app requires the use of two services, each running on a different system. One service has a 4% chance of failing at any given time and the other has a 10% chance of failing. Assume the network is always available. What is the approximate probability that the overall system will fail?

- (a) 0.4%
- (b) 7.0%
- (c) 3.6%
- (d) 13.6%

Question 3

The main challenge with *fail-restart* behavior is:

- (a) Network connections will have to be re-established.
- (b) The system may fail again.
- (c) The system may have stale state.
- (d) The network could get partitioned.

Question 4

The *transport layer* differs from the network layer because it:

- (a) Enables applications to communicate with each other.
- (b) Routes packets across multiple networks to their destination.
- (c) Provides a way to send structured data over the network.
- (d) Makes multiple connected physical networks look like a single virtual network.

Question 5

Why does the Network Time Protocol (NTP) use UDP instead of TCP?

- (a) To avoid added latency due to retransmission.
- (b) In-order delivery is not essential.
- (c) Minor data corruption is acceptable.
- (d) UDP operates at a lower layer in the network stack

Question 6

An RPC *client stub*:

- (a) Is a template that allows the programmer to insert code to call a remote procedure.
- (b) Is automatically generated code that the programmer calls to invoke the remote procedure.**
- (c) Contains a portable implementation of the remote procedure that can be run on another system.
- (d) Abstracts all the code needed to look up remote interfaces and establish a connection to the server.

Question 7

Idempotent functions can be run any number of times without harmful side effects. They are desirable in remote procedure calls because:

- (a) Their payloads are smaller than comparable non-idempotent functions.
- (b) They can execute more efficiently.
- (c) They simplify handling failure.
- (d) They support the use of TCP or UDP.

Question 8

The advantage of a *multi-canonical* data serialization (marshaling) format such as used in DCE RPC is:

- (a) The same message can be sent to systems with different hardware architectures or programming languages.
- (b) In most cases, at least one system will not have to convert data from its native representation.**
- (c) It allows entire data structures to be transmitted.
- (d) It stores data in a simple, human-readable format.

Question 9

Microsoft's DCOM used *pinging* in addition to remote reference counting because:

- (a) It enables servers to validate that the reference counts are valid.
- (b) Servers can load remote objects on demand from the client.
- (c) It provides a mechanism for servers to tell clients that they are alive.
- (d) It allows servers to clean up objects if a client or network dies.

Question 10

One advantage *protocol buffers* have over JSON is:

- (a) They are language-independent and platform-neutral.
- (b) They support explicit typing to identify the data in the message.
- (c) Data is stored in a compact binary format and is more efficient to parse.
- (d) They can serialize structured data.

Question 11

Because it is designed for web services, gRPC does not support:

- (a) Stub functions generation.
- (b) An interface definition language.
- (c) Distributed garbage collection.
- (d) Timeouts for a remote procedure call

Question 12

A drift compensation function

- (a) Keeps the time from suddenly jumping forward or backward when synchronizing.
- (b) Is applied to try to minimize the effects of jitter on the local clock.
- (c) Provides an estimate of how much the local clock is drifting relative to a known good clock.
- (d) Reduces the effect of network delays when synchronizing clocks.

Question 13

The Berkeley clock synchronization algorithm differs from Cristian's algorithm because:

- (a) It assumes jitter does not exist.
- (b) It assumes that no machine has an accurate time source.**
- (c) It applies a drift compensation function to keep clocks accurate.
- (d) It assumes all machines are on the same local area network.

Question 14

The following activity takes place:

Client sends a request at [client's time] = 44

Server receives the request at [server's time] = 32

Server creates a response with a timestamp = 35

Server sends the response at [server's time] = 36

Client receives the response at [client's time] = 54

Using **Cristian's algorithm**, to what value does the client set its clock?

$$\begin{aligned}\text{New time} &= T_{\text{server}} + \frac{1}{2}(54-44) \\ &= T_{\text{server}} + \frac{1}{2}(10) = T_{\text{server}} + 5 \\ &= 35 + 5 = 40\end{aligned}$$

Question 15

If the best-case round-trip network transit time is 4, what is the computed error of this synchronization? Express your answer in the form $\pm N$.

- Measured round-trip time (RTT) = 10
- Best-case RTT = 4
- Uncertainty = $10 - 4 = 6$
- Maximum error = ± 3

Question 16

Causally ordered multicasts imply

- (a) Single-source FIFO ordering.
- (b) Total ordering.
- (c) Global time ordering.
- (d) None of the above.

Question 17

A *precedence vector* is used to allow a process in a group to:

- (a) Decide whether two messages are concurrent.
- (b) Agree with other processes on the sequence of concurrent events.
- (c) Find out whether it needs to wait for other messages to arrive before delivering the received message.
- (d) Determine the path that a message took to arrive at this process.


Question 18

The diagram on the right represents three processes that communicate with each other. Each labeled dot represents an event. Arrows represent messages between processes. The horizontal axis represents time.

What is the Lamport timestamp of event h if all clocks are initially set to 0 (i.e., event a gets a Lamport timestamp of 1)?


Question 18

What is the Lamport timestamp of event h if all clocks are initially set to 0 (i.e., event a gets a Lamport timestamp of 1)?




Question 19

In the same diagram, using Lamport's happened-before definition, list all the events that happened before event j .



Question 20

Using the same diagram, vector timestamps are assigned to each event in the form $\langle P_0, P_1, P_2 \rangle$, with initial clock values of $\langle 0, 0, 0 \rangle$ at each process. What is the vector clock of event j ?



Question 21

PIM-DM (Protocol Independent Multicast - Dense Mode) doesn't scale well compared to PIM-SM (Sparse Mode) because:

- (a) It requires setting up rendezvous points and configuring each router to know about them.
- (b) It first floods the entire network with multicast traffic before routers prune away some paths.**
- (c) Packets have a hop count limit to avoid routing loops, which limits their reach.
- (d) It is easy for ISPs to block multicast packets.

Question 22

IP multicast provides this form of multicast delivery:

- (a) Totally-ordered atomic.
- (b) Totally-ordered reliable.
- (c) Causally-ordered reliable.
- (d) Unordered unreliable.

Question 23

The two armies problem illustrates that:

- (a) Consensus can be reached if one system acknowledges each message from the other system.
- (b) It is impossible to reach consensus between two systems connected by an unreliable communication channel.
- (c) Reaching consensus between two systems requires multiple back-and-forth rounds of acknowledgments.
- (d) Consensus can be achieved by having an intermediate system coordinate the messages from both systems.

Question 24

Virtual synchrony implements:

- (a) Atomic multicasts.
- (b) Causally-ordered reliable multicasts.
- (c) Partially-ordered reliable multicasts.
- (d) Sync-ordered unreliable multicasts.

Question 25

In virtual synchrony, each group member must hold on to a message until:

- (a) A view change is initiated.
- (b) All messages that causally precede it have been received.
- (c) It sends an acknowledgment to the sender.
- (d) It receives confirmation that the entire group received it.

The End

The end

Exam 1 Grade Distribution

Average	70.9
σ	18.4
Highest	100
Lowest	32

Top 10%	≥ 96
Top 20%	≥ 88
Bottom 10%	≤ 44
Bottom 20%	≤ 56

Approximate grades

A	≥ 82
B+	≥ 73
B	≥ 64
C+	≥ 54
C	≥ 45

