Confluent CCAAK Practice Questions

Shared by Diaz on 10-11-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A company is setting up a log ingestion use case where they will consume logs from numerous systems. The company wants to tune Kafka for the utmost throughput.

In this scenario, what acknowledgment setting makes the most sense?

Options:

A- acks=0
B- acks=1
C- acks=all
D- acks=undefined

Answer:

A

Explanation:

acks=0 provides the highest throughput because the producer does not wait for any acknowledgment from the broker. This minimizes latency and maximizes performance.

However, it comes at the cost of no durability guarantees --- messages may be lost if the broker fails before writing them. This setting is suitable when throughput is critical and occasional data loss is acceptable, such as in some log ingestion use cases where logs are also stored elsewhere.

Question 2

Question Type: MultipleChoice

Which connector type takes data from a topic and sends it to an external data system?

Options:

A- Sink Connector
B- Source Connector
C- Streams Connector

D- SvsIog Connector

Answer:

A

Explanation:

A Sink Connector reads data from a Kafka topic and writes it to an external data system, such as a database, file system, or cloud service.

P2P exams

Question 3

Question Type: MultipleChoice

Your organization has a mission-critical Kafka cluster that must be highly available. A Disaster Recovery (DR) cluster has been set up using Replicator, and data is continuously being replicated from source cluster to the DR cluster. However, you notice that the message on offset 1002 on source cluster does not seem to match with offset 1002 on the destination DR cluster.

Which statement is correct?

Options:

A- The DR cluster is lagging behind updates; once the DR cluster catches up, the messages will match.
B- The message on DR cluster got over-written accidently by another application.
C- The offsets for the messages on the source, destination cluster may not match.
D- The message was updated on source cluster, but the update did not flow into destination DR cluster and errored.

P2P exams

Answer:

C

Explanation:

When using Confluent Replicator (or MirrorMaker), offsets are not preserved between the source and destination Kafka clusters. Messages are replicated based on content, but they are assigned new offsets in the DR (destination) cluster. Therefore, offset 1002 on the source and offset 1002 on the DR cluster likely refer to different messages, which is expected behavior.

Question 4

Question Type: MultipleChoice

A company is setting up a log ingestion use case where they will consume logs from numerous systems. The company wants to tune Kafka for the utmost throughput.

In this scenario, what acknowledgment setting makes the most sense?

Options:

A- acks=0
B- acks=1
C- acks=all
D- acks=undefined

Answer:

A

Explanation:

acks=0 provides the highest throughput because the producer does not wait for any acknowledgment from the broker. This minimizes latency and maximizes performance.

However, it comes at the cost of no durability guarantees --- messages may be lost if the broker fails before writing them. This setting is suitable when throughput is critical and occasional data loss is acceptable, such as in some log ingestion use cases where logs are also stored elsewhere.

Question 5

Question Type: MultipleChoice

Which secure communication is supported between the REST proxy and REST clients?

Options:

A- TLS (HTTPS)

B- MD5
C- SCRAM
D- Kerberos

Answer:

A

Question 6

Question Type: MultipleChoice

What are important factors in sizing a ksqlDB cluster? (Choose three.)

Options:

A- Data Schema
B- Number of Queries
C- Number of Partitions
D- Data Encryption
E- Topic Data Retention

Answer:

A, B, C

Explanation:

The complexity of the schema (number of fields, data types, etc.) affects processing and memory usage.

Each ksqlDB persistent query consumes resources (CPU, memory), so more queries require more capacity.

More partitions increase parallelism, but also resource usage, especially in scaling and state management.

Question 7

Question Type: MultipleChoice

A broker in the Kafka cluster is currently acting as the Controller.

Which statement is correct?

Options:

A- It can have topic partitions.
B- It is given precedence for replication to and from replica followers.
C- All consumers are allowed to fetch messages only from this server.
D- It is responsible for sending leader information to all producers.

Answer:

A

Explanation:

The Controller broker is a regular broker that also takes on additional responsibilities for managing cluster metadata, such as leader elections and partition assignments. It still hosts topic partitions and participates in replication like any other broker.

Question 8

Question Type: MultipleChoice

Your Kafka cluster has four brokers. The topic t1 on the cluster has two partitions, and it has a replication factor of three. You create a Consumer Group with four consumers, which subscribes to t1.

In the scenario above, how many Controllers are in the Kafka cluster?

Options:

A- One
B- two
C- three
D- Four

Answer:

A

Explanation:

In a Kafka cluster, only one broker acts as the Controller at any given time. The Controller is responsible for managing cluster metadata, such as partition leadership and broker status. Even if the cluster has multiple brokers (in this case, four), only one is elected as the Controller, and others serve as regular brokers. If the current Controller fails, another broker is automatically elected to take its place.

P2P exams

Question 9

Question Type: MultipleChoice

How can load balancing of Kafka clients across multiple brokers be accomplished?

Options:

A- Partitions
B- Replicas
C- Offsets
D- Connectors

Answer:

A

P2P exams

Explanation:

Partitions are the primary mechanism for achieving load balancing in Kafka. When a topic has multiple partitions, Kafka clients (producers and consumers) can distribute the load across brokers hosting these partitions.

Question 10

Question Type: MultipleChoice

You are managing a Kafka cluster with five brokers (broker id '0', '1','2','3','4') and three ZooKeepers. There are 100 topics, five partitions for each topic, and replication factor three on the cluster. Broker id '0' is currently the Controller, and this broker suddenly fails.

Which statements are correct? (Choose three.)

Options:

A- Kafka uses ZooKeeper's ephemeral node feature to elect a controller.
B- The Controller is responsible for electing Leaders among the partitions and replicas.
C- The Controller uses the epoch number to prevent a split brain scenario.
D- The broker id is used as the epoch number to prevent a split brain scenario.
E- The number of Controllers should always be equal to the number of brokers alive in the cluster.
F- The Controller is responsible for reassigning partitions to the consumers in a Consumer Group.

Answer:

A, B, C

Explanation:

Kafka relies on ZooKeeper's ephemeral nodes to detect if a broker (controller) goes down and to elect a new controller.

The controller manages partition leadership assignments and handles leader election when a broker fails.

The epoch number ensures coordination and avoids outdated controllers acting on stale data.

P2P exams

Question 11

Question Type: MultipleChoice

In certain scenarios, it is necessary to weigh the trade-off between latency and throughput. One method to increase throughput is to configure batching of messages.

In addition to batch.size, what other producer property can be used to accomplish this?

Options:

A- sendbufferbytes
B- linger.ms
C- compression
D- delivery.timeout.ms

Answer:

B

Explanation:

The linger.ms property sets the maximum time the producer will wait before sending a batch, even if the batch isn't full. By increasing linger.ms, you allow more records to accumulate in the batch.

Question 12

Question Type: MultipleChoice

If a broker's JVM garbage collection takes too long, what can occur?

Options:

A- There will be a trigger of the broker's log cleaner thread.
B- ZooKeeper believes the broker to be dead.
C- There is backpressure to, and pausing of, Kafka clients.
D- Log files written to disk are loaded into the page cache.

Answer:

B

Explanation:

If the broker's JVM garbage collection (GC) pause is too long, it may fail to send heartbeats to ZooKeeper within the expected interval. As a result, ZooKeeper considers the broker dead, and the broker may be removed from the cluster, triggering leader elections and partition reassignments.

Question 13

Question Type: MultipleChoice

When using Kafka ACLs, when is the resource authorization checked?

Options:

A- Each time the resource is accessed.
B- The initial time the resource is accessed.
C- Each time the resource is accessed within the configured authorization interval.
D- When the client connection is first established.

Answer:

A

Explanation:

Kafka ACLs (Access Control Lists) perform authorization checks every time a client attempts to access a resource (e.g., topic, consumer group). This ensures continuous enforcement of permissions, not just at connection time or intervals. This approach provides fine-grained security, preventing unauthorized actions at any time during a session.

To Get Premium Files for CCAAK Visit

https://www.p2pexams.com/products/ccaak

For More Free Questions Visit

https://www.p2pexams.com/confluent/pdf/ccaak

