Download Confluent CCAAK Exam Dumps Free

Shared by Guerrero on 17-06-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

What does Kafka replication factor provide? (Choose two.)

Options:

A- Performance
B- Durability
C- Availability
D- Security

P2P exams

Answer:

B, C

Explanation:

Replication ensures that multiple copies of data exist across different brokers, so data is not lost if one broker fails.

With multiple replicas, Kafka can continue to serve data even if the leader or one replica fails, maintaining service availability.

Question 2

Question Type: MultipleChoice

P2P exams

Which option is a valid Kafka Topic cleanup policy? (Choose two.)

Options:

A- delete
B- default
C- compact
D- cleanup

Answer:

A, C

Explanation:

The delete policy deletes old log segments when they exceed the retention period or size.

The compact policy retains only the latest record for each key, enabling efficient key-based storage.

Question 3

Question Type: MultipleChoice

What is the atomic unit of data in Kafka?

Options:

A- Partition
B- Message
C- Topic
D- Offset

Answer:

B

Explanation:

The message (also called a record) is the atomic unit of data in Kafka. It consists of a key, value, and optional headers and metadata. Messages are written to partitions and read by consumers.

Question 4

Question Type: MultipleChoice

You are using Confluent Schema Registry to provide a RESTful interface for storing and retrieving

schemas.

Which types of schemas are supported? (Choose three.)

Options:

A- Avro
B- gRPC
C- JSON
D- Thrift
E- Protobuf

Answer:

A, C, E

Explanation:

Avro is the original and most commonly used schema format supported by Schema Registry.

Confluent Schema Registry supports JSON Schema for validation and compatibility checks.

Protocol Buffers (Protobuf) are supported for schema management in Schema Registry.

Question 5

Question Type: MultipleChoice

If the Controller detects the failure of a broker that was the leader for some partitions, which actions will be taken? (Choose two.)

Options:

A- The Controller waits for a new leader to be nominated by ZooKeeper.
B- The Controller persists the new leader and ISR list to ZooKeeper.
C- The Controller sends the new leader and ISR list changes to all brokers.
D- The Controller sends the new leader and ISR list changes to all producers and consumers.

Answer:

B, C

Explanation:

The Controller updates ZooKeeper with the new leader and in-sync replica (ISR) information to maintain metadata consistency.

Brokers need this information to correctly route client requests and continue replication.

Question 6

Question Type: MultipleChoice

You have a Kafka cluster with topics t1 and t2. In the output below, topic t2 shows Partition 1 with a leader ''-1''.

What is the most likely reason for this?

...

$ kafka-topics --zookeeper localhost:2181 --describe --topic t1

Topic:t1 PartitionCount 1 ReplicationFactor 1 Configs:

Topic: t1 Partition: 0 Leader: 0 Replicas: 0 Isr: 0

$ kafka-topics --zookeeper localhost:2181 --describe --topic t2

Topic:t2 PartitionCount 2 ReplicationFactor 1 Configs:

Topic: t2 Partition: 0 Leader: 0 Replicas: 0 Isr: 0

Topic: t2 Partition: 1 Leader: -1 Replicas: 1 Isr:

Options:

A- Broker 1 failed.
B- Leader shows ''-1'' while the log cleaner thread runs on Broker 1.
C- Compression has been enabled on Broker 1.
D- Broker 1 has another partition clashing with the same name.

Answer:

A

Explanation:

A Leader of -1 indicates that no broker is currently the leader for that partition. This usually happens when the only replica for that partition is unavailable, often due to the associated broker (in this case, Broker 1) failing or being offline. Kafka cannot elect a leader if no replica is in the in-sync replica (ISR) list, which leads to leader = -1.

Question 7

Question Type: MultipleChoice

What is the relationship between topics and partitions? (Select two.)

Options:

A- Atopic always has one partition.
B- Atopic may have more than one partition.
C- A partition is always linked to a single topic.
D- A partition may have more than one topic.
E- There is no relationship between topics and partitions.

Answer:

B, C

Explanation:

Kafka topics are split into one or more partitions to enable parallelism and scalability.

Each partition belongs to exactly one topic; it cannot span multiple topics.

Question 8

Question Type: MultipleChoice

Which ksqlDB statement produces data that is persisted into a Kafka topic?

Options:

A- SELECT (Pull Query)
B- SELECT (Push Query)
C- INSERT VALUES
D- CREATE TABLE

Answer:

C

Explanation:

INSERT VALUES is used to write data directly into a Kafka topic through a ksqlDB stream or table. This data is persisted.

Question 9

Question Type: MultipleChoice

A company has an existing Kafka cluster running without SSL/TLS enabled. The customer wants to enable SSL on brokers to secure data in transit, but they would like to give applications connecting to this cluster some time to migrate to using SSL connection instead of putting a hard stop.

Which solution will meet the customer's requirements?

Options:

A- Enable SSL on the current Listener, and do not enable mTLS.
B- Modify the advertised listeners setting on brokers to use SSL.
C- Create a new listener with SSL enabled.
D- Enable SSL on the current listener, and do not implement SSL on application side.

Answer:

C

Explanation:

Kafka supports multiple listeners, allowing you to run PLAINTEXT and SSL simultaneously. By creating a new SSL-enabled listener (e.g., on a different port), existing applications can continue using PLAINTEXT while gradually migrating to the SSL listener. This approach avoids downtime and gives clients time to adapt without enforcing a hard cutover.

Question 10

Question Type: MultipleChoice

Kafka Connect is running on a two node cluster in distributed mode. The connector is a source connector that pulls data from Postgres tables (users/payment/orders), writes to topics with two partitions, and with replication factor two. The development team notices that the data is lagging behind.

What should be done to reduce the data lag*?

The Connector definition is listed below:

{

"name": "confluent-postgresql-source",

"connector class": "PostgresSource",

"topic.prefix": "postgresql_",

& nbsp;& nbsp;& nbsp;...

"db.name": "postgres",

"table.whitelist": "users.payment.orders'',

"timestamp.column.name": "created_at",

"output.data format": "JSON",

"db.timezone": "UTC",

"tasks.max": "1"

}

Options:

A- Increase the number of Connect Nodes.

B- Increase the number of Connect Tasks (tasks max value).

C- Increase the number of partitions.
D- Increase the replication factor and increase the number of Connect Tasks.

Answer:

B

Explanation:

The connector is currently configured with 'tasks.max': '1', which means only one task is handling all tables (users, payment, orders). This can create a bottleneck and lead to lag. Increasing tasks.max allows Kafka Connect to parallelize work across multiple tasks, which can pull data from different tables concurrently and reduce lag.

P2P exams

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