Confluent CCDAK Practice Questions

Shared by Frederick on 10-11-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

To import data from external databases, I should use

Options:

A- Confluent REST Proxy
B- Kafka Connect Sink
C- Kafka Streams
D- Kafka Connect Source

Answer:

D

Explanation:

Kafka Connect Sink is used to export data from Kafka to external databases and Kafka Connect Source is used to import from external databases into Kafka.

Question 2

Question Type: MultipleChoice

When using plain JSON data with Connect, you see the following error messageorg.apache.kafka.connect.errors.DataExceptionJsonDeserializer with schemas.enable requires "schema" and "payload" fields and may not contain additional fields. How will you fix the error?

Options:

A- Set key.converter, value.converter to JsonConverter and the schema registry url
B- Use Single Message Transforms to add schema and payload fields in the message
C- Set key.converter.schemas.enable and value.converter.schemas.enable to false
D- Set key.converter, value.converter to AvroConverter and the schema registry url

Answer:

C

Explanation:

You will need to set the schemas.enable parameters for the converter to false for plain text with no schema.

Question 3

Question Type: MultipleChoice

What is not a valid authentication mechanism in Kafka?

Options:

A- SASL/GSSAPI
B- SASL/SCRAM
C- SAML
D- SSL

Answer:

C

Explanation:

Learn more about security herehttps://kafka.apache.org/documentation/#security

Question 4

Question Type: MultipleChoice

How will you find out all the partitions where one or more of the replicas for the partition are not in-sync with the leader?

Options:

A- kafka-topics.sh --bootstrap-server localhost:9092 --describe --unavailable- partitions
B- kafka-topics.sh --zookeeper localhost:2181 --describe --unavailable- partitions
C- kafka-topics.sh --broker-list localhost:9092 --describe --under-replicated-partitions
D- kafka-topics.sh --zookeeper localhost:2181 --describe --under-replicated-partitions

Answer:

D

Question 5

Question Type: MultipleChoice

What happens if you write the following code in your producer?
producer.send(producerRecord).get()

Options:

A- Compression will be increased
B- Throughput will be decreased
C- It will force all brokers in Kafka to acknowledge the producerRecord
D- Batching will be increased

Answer:

B

Explanation:

Using Future.get() to wait for a reply from Kafka will limit throughput.

Question 6

Question Type: MultipleChoice

You have a consumer group of 12 consumers and when a consumer gets killed by the process

management system, rather abruptly, it does not trigger a graceful shutdown of your consumer. Therefore, it takes up to 10 seconds for a rebalance to happen. The business would like to have a 3 seconds rebalance time. What should you do? (select two)

Options:

A- Increase session.timeout.ms
B- Decrease session.timeout.ms
C- Increase heartbeat.interval.ms
D- decrease max.poll.interval.ms
E- increase max.poll.interval.ms
F- Decrease heartbeat.interval.ms

P2P exams

Answer:

B, E

Explanation:

session.timeout.ms must be decreased to 3 seconds to allow for a faster rebalance, and the heartbeat thread must be quicker, so we also need to decrease heartbeat.interval.ms

Question 7

Question Type: MultipleChoice

A consumer application is using KafkaAvroDeserializer to deserialize Avro messages. What happens if message schema is not present in AvroDeserializer local cache?

P2P exams

Options:

A- Throws SerializationException
B- Fails silently
C- Throws DeserializationException
D- Fetches schema from Schema Registry

Answer:

D

Explanation:

First local cache is checked for the message schema. In case of cache miss, schema is pulled from the schema registry. An exception will be thrown in the Schema Registry does not have the schema (which should never happen if you set it up properly)

Question 8

Question Type: MultipleChoice

To produce data to a topic, a producer must provide the Kafka client with...

Options:

A- the list of brokers that have the data, the topic name and the partitions list
B- any broker from the cluster and the topic name and the partitions list
C- all the brokers from the cluster and the topic name
D- any broker from the cluster and the topic name

Answer:

D

Explanation:

All brokers can respond to a Metadata request, so a client can connect to any broker in the cluster and then figure out on its own which brokers to send data to.

Question 9

Question Type: MultipleChoice

What is a generic unique id that I can use for messages I receive from a consumer?

Options:

A- topic + partition + timestamp
B- topic + partition + offset
C- topic + timestamp

Answer:

B

Explanation:

(Topic,Partition,Offset) uniquely identifies a message in Kafka

P2P exams

Question 10

Question Type: MultipleChoice

A consumer starts and has auto.offset.reset=none, and the topic partition currently has data for offsets going from 45 to 2311. The consumer group has committed the offset 10 for the topic before. Where will the consumer read from?

Options:

A- offset 45
B- offset 10
C- it will crash
D- offset 2311

P2P exams

Answer:

C

Explanation:

auto.offset.reset=none means that the consumer will crash if the offsets it's recovering from have been deleted from Kafka, which is the case here, as 10 < 45

Question 11

Question Type: MultipleChoice

You are doing complex calculations using a machine learning framework on records fetched from a Kafka topic. It takes more about 6 minutes to process a record batch, and the consumer enters rebalances even though it's still running. How can you improve this scenario?

Options:

A- Increase max.poll.interval.ms to 600000
B- Increase heartbeat.interval.ms to 600000
C- Increase session.timeout.ms to 600000
D- Add consumers to the consumer group and kill them right away

Answer:

A

Explanation:

Here, we need to change the setting max.poll.interval.ms (default 300000) to its double in order to tell Kafka a consumer should be considered dead if the consumer only if it hasn't called the .poll() method in 10 minutes instead of 5.

Question 12

Question Type: MultipleChoice

What happens when broker.rack configuration is provided in broker configuration in Kafka cluster?

Options:

A- You can use the same broker.id as long as they have different broker.rack configuration
B- Replicas for a partition are placed in the same rack
C- Replicas for a partition are spread across different racks
D- Each rack contains all the topics and partitions, effectively making Kafka highly available

Answer:

C

Explanation:

Partitions for newly created topics are assigned in a rack alternating manner, this is the only change broker.rack does

P2P exams

P2P exams

To Get Premium Files for CCDAK Visit

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

For More Free Questions Visit

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

