



CCDAK^{Q&As}

Confluent Certified Developer for Apache Kafka Certification
Examination

Pass Confluent CCDAK Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/ccdak.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Confluent
Official Exam Center

- ⚙ **Instant Download** After Purchase
- ⚙ **100% Money Back** Guarantee
- ⚙ **365 Days** Free Update
- ⚙ **800,000+** Satisfied Customers



**QUESTION 1**

Suppose you have 6 brokers and you decide to create a topic with 10 partitions and a replication factor of 3. The brokers 0 and 1 are on rack A, the brokers 2 and 3 are on rack B, and the brokers 4 and 5 are on rack C. If the leader for partition 0 is on broker 4, and the first replica is on broker 2, which broker can host the last replica? (select two)

- A. 6
- B. 1
- C. 2
- D. 5
- E. 0
- F. 3

Correct Answer: BE

When you create a new topic, partitions replicas are spread across racks to maintain availability. Hence, the Rack A, which currently does not hold the topic partition, will be selected for the last replica

QUESTION 2

You are using JDBC source connector to copy data from 3 tables to three Kafka topics. There is one connector created with max.tasks equal to 2 deployed on a cluster of 3 workers. How many tasks are launched?

- A. 2
- B. 1
- C. 3
- D. 6

Correct Answer: A

here, we have three tables, but the max.tasks is 2, so that's the maximum number of tasks that will be created

QUESTION 3

A client connects to a broker in the cluster and sends a fetch request for a partition in a topic. It gets an exception Not Leader For Partition Exception in the response. How does client handle this situation?

- A. Get the Broker id from Zookeeper that is hosting the leader replica and send request to it
- B. Send metadata request to the same broker for the topic and select the broker hosting the leader replica
- C. Send metadata request to Zookeeper for the topic and select the broker hosting the leader replica



D. Send fetch request to each Broker in the cluster

Correct Answer: B

In case the consumer has the wrong leader of a partition, it will issue a metadata request. The Metadata request can be handled by any node, so clients know afterwards which broker are the designated leader for the topic partitions. Produce and consume requests can only be sent to the node hosting partition leader.

QUESTION 4

To prevent network-induced duplicates when producing to Kafka, I should use

A. `max.in.flight.requests.per.connection=1`

B. `enable.idempotence=true`

C. `retries=200000`

D. `batch.size=1`

Correct Answer: B

Producer idempotence helps prevent the network introduced duplicates. More details here <https://cwiki.apache.org/confluence/display/KAFKA/Idempotent+Producer>

QUESTION 5

We have a store selling shoes. What dataset is a great candidate to be modeled as a KTable in Kafka Streams?

A. Money made until now

B. The transaction stream

C. Items returned

D. Inventory contents right now

Correct Answer: C

QUESTION 6

There are 3 producers writing to a topic with 5 partitions. There are 5 consumers consuming from the topic. How many Controllers will be present in the cluster?

A. 3

B. 5

C. 2



D. 1

Correct Answer: D

There is only one controller in a cluster at all times.

QUESTION 7

```
while (true) {  
  
    ConsumerRecords records = consumer.poll(100); try {  
  
        consumer.commitSync();  
  
    } catch (CommitFailedException e) {  
  
        log.error("commit failed", e)  
  
    }  
  
    for (ConsumerRecord record records)  
  
    {  
  
        System.out.printf("topic = %s, partition = %s, offset =  
  
        %d, customer = %s, country = %s  
  
        ",  
  
        record.topic(), record.partition(),  
  
        record.offset(), record.key(), record.value());  
  
    }  
  
}
```

What kind of delivery guarantee this consumer offers?

- A. Exactly-once
- B. At-least-once
- C. At-most-once

Correct Answer: C

Here offset is committed before processing the message. If consumer crashes before processing the message, message will be lost when it comes back up.

QUESTION 8



A consumer starts and has `auto.offset.reset=latest`, and the topic partition currently has data for offsets going from 45 to 2311. The consumer group has committed the offset 643 for the topic before. Where will the consumer read from?

- A. it will crash
- B. offset 2311
- C. offset 643
- D. offset 45

Correct Answer: C

The offsets are already committed for this consumer group and topic partition, so the property `auto.offset.reset` is ignored

QUESTION 9

What data format isn't natively available with the Confluent REST Proxy?

- A. avro
- B. binary
- C. protobuf
- D. json

Correct Answer: C

Protocol buffers isn't a natively supported type for the Confluent REST Proxy, but you may use the binary format instead

QUESTION 10

What Java library is KSQL based on?

- A. Kafka Streams
- B. REST Proxy
- C. Schema Registry
- D. Kafka Connect

Correct Answer: A

KSQL is based on Kafka Streams and allows you to express transformations in the SQL language that get automatically converted to a Kafka Streams program in the backend

QUESTION 11



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

- 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

Correct Answer: C

QUESTION 12

You are sending messages with keys to a topic. To increase throughput, you decide to increase the number of partitions of the topic. Select all that apply.

- A. All the existing records will get rebalanced among the partitions to balance load
- B. New records with the same key will get written to the partition where old records with that key were written
- C. New records may get written to a different partition
- D. Old records will stay in their partitions

Correct Answer: CD

Increasing the number of partition causes new messages keys to get hashed differently, and breaks the guarantee "same keys goes to the same partition". Kafka logs are immutable and the previous messages are not re-shuffled

QUESTION 13

A consumer is configured with enable.auto.commit=false. What happens when close() is called on the consumer object?

- A. The uncommitted offsets are committed
- B. A rebalance in the consumer group will happen immediately
- C. The group coordinator will discover that the consumer stopped sending heartbeats. It will cause rebalance after session.timeout.ms

Correct Answer: B

Calling close() on consumer immediately triggers a partition rebalance as the consumer will not be available anymore.

QUESTION 14

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?

- 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

Correct Answer: C

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

QUESTION 15

A topic has three replicas and you set min.insync.replicas to 2. If two out of three replicas are not available, what happens when a consume request is sent to broker?

- A. Data will be returned from the remaining in-sync replica
- B. An empty message will be returned
- C. NotEnoughReplicasException will be returned
- D. A new leader for the partition will be elected

Correct Answer: A

With this configuration, a single in-sync replica is still readable, but not writeable if the producer using acks=all

[Latest CCDAK Dumps](#)

[CCDAK VCE Dumps](#)

[CCDAK Braindumps](#)