

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.lead4pass.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

To read data from a topic, the following configuration is needed for the consumers

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

Correct Answer: B

All brokers can respond to Metadata request, so a client can connect to any broker in the cluster.

QUESTION 2

An ecommerce website sells some custom made goods. What's the natural way of modeling this data in Kafka streams?

- A. Purchase as stream, Product as stream, Customer as stream
- B. Purchase as stream, Product as table, Customer as table
- C. Purchase as table, Product as table, Customer as table
- D. Purchase as stream, Product as table, Customer as stream

Correct Answer: B

Mostly-static data is modeled as a table whereas business transactions should be modeled as a stream.

QUESTION 3

If I supply the setting `compression.type=snappy` to my producer, what will happen? (select two) A. The Kafka brokers have to de-compress the data

- B. The Kafka brokers have to compress the data
- C. The Consumers have to de-compress the data
- D. The Consumers have to compress the data
- E. The Producers have to compress the data

Correct Answer: C

Kafka transfers data with zero copy and no transformation. Any transformation (including compression) is the responsibility of clients.

QUESTION 4

What isn't a feature of the Confluent schema registry?

- A. Store avro data
- B. Enforce compatibility rules
- C. Store schemas

Correct Answer: A

Data is stored on brokers.

QUESTION 5

There are 3 brokers in the cluster. You want to create a topic with a single partition that is resilient to one broker failure and one broker maintenance. What is the replication factor will you specify while creating the topic?

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

Correct Answer: B

1 is not possible as it doesn't provide resilience to failure, 2 is not enough as if we take a broker down for maintenance, we cannot tolerate a broker failure, and 6 is impossible as we only have 3 brokers (RF cannot be greater than the number of brokers). Here the correct answer is 3

QUESTION 6

How do you create a topic named test with 3 partitions and 3 replicas using the Kafka CLI?

- A. `bin/kafka-topics.sh --create --broker-list localhost:9092 --replication-factor 3 --partitions 3 --topic test`
- B. `bin/kafka-topics-create.sh --zookeeper localhost:9092 --replication-factor 3 --partitions 3 --topic test`
- C. `bin/kafka-topics.sh --create --bootstrap-server localhost:9092 --replication-factor 3 --partitions 3 --topic test`
- D. `bin/kafka-topics.sh --create --bootstrap-server localhost:2181 --replication-factor 3 --partitions 3 --topic test`

Correct Answer: C

As of Kafka 2.3, the `kafka-topics.sh` command can take `--bootstrap-server localhost:9092` as an argument. You could also use the (now deprecated) option of `--zookeeper localhost:2181`.

QUESTION 7

If I produce to a topic that does not exist, and the broker setting `auto.create.topic.enable=true`, what will happen?

- A. Kafka will automatically create the topic with 1 partition and 1 replication factor
- B. Kafka will automatically create the topic with the indicated producer settings `num.partitions` and `default.replication.factor`
- C. Kafka will automatically create the topic with the broker settings `num.partitions` and `default.replication.factor`
- D. Kafka will automatically create the topic with `num.partitions=#of brokers` and `replication.factor=3`

Correct Answer: C

The broker settings comes into play when a topic is auto created

QUESTION 8

Partition leader election is done by

- A. The consumers
- B. The Kafka Broker that is the Controller
- C. Zookeeper
- D. Vote amongst the brokers

Correct Answer: C

The Controller is a broker that is responsible for electing partition leaders

QUESTION 9

Which of the following is true regarding thread safety in the Java Kafka Clients?

- A. One Producer can be safely used in multiple threads
- B. One Consumer can be safely used in multiple threads
- C. One Consumer needs to run in one thread
- D. One Producer needs to be run in one thread

Correct Answer: AC

KafkaConsumer is not thread-safe, KafkaProducer is thread safe.

QUESTION 10

Kafka is configured with following parameters - log.retention.hours = 168 log.retention.minutes = 168 log.retention.ms = 168 How long will the messages be retained for?

- A. Broker will not start due to bad configuration
- B. 168 ms
- C. 168 hours
- D. 168 minutes

Correct Answer: B

If more than one similar config is specified, the smaller unit size will take precedence.

QUESTION 11

```
StreamsBuilder builder = new StreamsBuilder();  
  
KStream textLines = builder.stream("word-count-input");  
  
KTable wordCounts = textLines  
    .mapValues(textLine -> textLine.toLowerCase())  
    .flatMapValues(textLine -> Arrays.asList(textLine.split("\\W+"))) .selectKey((key, word) -> word)  
    .groupByKey()  
    .count(Materialized.as("Counts"));  
  
wordCounts.toStream().to("word-count-output", Produced.with(Serdes.String(), Serdes.Long()));  
  
builder.build();
```

What is an adequate topic configuration for the topic word-count-output?

- A. max.message.bytes=10000000
- B. cleanup.policy=delete
- C. compression.type=lz4
- D. cleanup.policy=compact

Correct Answer: D

Result is aggregated into a table with key as the unique word and value its frequency. We have to enable log compaction for this topic to align the topic's cleanup policy with KTable semantics.

QUESTION 12

What is returned by a producer.send() call in the Java API?

- A. Future object
- B. A Boolean indicating if the call succeeded
- C. Future object
- D. Unit

Correct Answer: C

See <https://kafka.apache.org/21/javadoc/org/apache/kafka/clients/producer/KafkaProducer.html>

QUESTION 13

A Zookeeper ensemble contains 5 servers. What is the maximum number of servers that can go missing and the ensemble still run?

- A. 3
- B. 4
- C. 2
- D. 1

Correct Answer: C

majority consists of 3 zk nodes for 5 nodes zk cluster, so 2 can fail

QUESTION 14

A consumer sends a request to commit offset 2000. There is a temporary communication problem, so the broker never gets the request and therefore never responds. Meanwhile, the consumer processed another batch and successfully committed offset 3000. What should you do?

- A. Add a new consumer to the group
- B. Use the kafka-consumer-group command to manually commit the offsets 2000 for the consumer group
- C. Restart the consumer
- D. Nothing

Correct Answer: D

In this case, because the offset 3000 has been committed and all the messages between 0 and 3000 have all been processed, it is okay not to have committed offset 2000. The right answer is to do "nothing", this behaviour is acceptable

QUESTION 15

The Controller is a broker that is... (select two)

- A. elected by Zookeeper ensemble
- B. is responsible for partition leader election
- C. elected by broker majority
- D. is responsible for consumer group rebalances

Correct Answer: AB

Controller is a broker that in addition to usual broker functions is responsible for partition leader election. The election of that broker happens thanks to Zookeeper and at any time only one broker can be a controller

[CCDAK VCE Dumps](#)

[CCDAK Study Guide](#)

[CCDAK Exam Questions](#)