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

You have a Kafka cluster and all the topics have a replication factor of 3. One intern at your company

stopped a broker, and accidentally deleted all the data of that broker on the disk.

What will happen if the broker is restarted?

A. The broker will start, and other topics will also be deleted as the broker data on the disk got deleted

B. The broker will start, and won\\'t be online until all the data it needs to have is replicated from other leaders

C. The broker will crash

D. The broker will start, and won\\'t have any data. If the broker comes leader, we have a data loss

Correct Answer: B

Kafka replication mechanism makes it resilient to the scenarios where the broker lose data on disk, but can recover from replicating from other brokers. This makes Kafka amazing!

QUESTION 2

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 3

Which of the following event processing application is stateless? (select two)

A. Read events from a stream and modifies them from JSON to Avro

B. Publish the top 10 stocks each day

C. Read log messages from a stream and writes ERROR events into a high-priority stream and the rest of the events into a low-priority stream

D. Find the minimum and maximum stock prices for each day of trading

Correct Answer: AC

Stateless means processing of each message depends only on the message, so converting from JSON to Avro or filtering a stream are both stateless operations

QUESTION 4

We would like to be in an at-most once consuming scenario. Which offset commit strategy would you recommend?

A. Commit the offsets on disk, after processing the data

B. Do not commit any offsets and read from beginning

C. Commit the offsets in Kafka, after processing the data

D. Commit the offsets in Kafka, before processing the data

Correct Answer: D

Here, we must commit the offsets right after receiving a batch from a call to .poll()

QUESTION 5

Compaction is enabled for a topic in Kafka by setting log.cleanup.policy=compact. What is true about log compaction?

A. After cleanup, only one message per key is retained with the first value

B. Each message stored in the topic is compressed

C. Kafka automatically de-duplicates incoming messages based on key hashes

D. After cleanup, only one message per key is retained with the latest value Compaction changes the offset of messages

Correct Answer: D

Log compaction retains at least the last known value for each record key for a single topic partition. All compacted log offsets remain valid, even if record at offset has been compacted away as a consumer will get the next highest offset.