DUMPSQUEEN

Confluent Certified Developer for Apache Kafka Certification Examination

Confluent CCDAK

Version Demo

Total Demo Questions: 10

Total Premium Questions: 150

Buy Premium PDF

https://dumpsqueen.com

support@dumpsqueen.com

QUESTION NO: 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

ANSWER: B

Explanation:

:

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 NO: 2

In Kafka, every broker... (select three)

A. contains all the topics and all the partitions

B. knows all the metadata for all topics and partitions

C. is a controller

D. knows the metadata for the topics and partitions it has on its disk

E. is a bootstrap broker

F. contains only a subset of the topics and the partitions

ANSWER: B E F

Explanation:

:

Kafka topics are divided into partitions and spread across brokers. Each brokers knows about all the metadata and each broker is a bootstrap broker, but only one of them is elected controller

QUESTION NO: 3

Which of the following errors are retriable from a producer perspective? (select two)

A. MESSAGE_TOO_LARGE

B. INVALID_REQUIRED_ACKS

C. NOT_ENOUGH_REPLICAS

D. NOT_LEADER_FOR_PARTITION

E. TOPIC_AUTHORIZATION_FAILED

ANSWER: C D

Explanation:

:

Both of these are retriable errors, others non-retriable errors. See the full list of errors and their "retriable" status herehttps://kafka.apache.org/protocol#protocol_error_codes

QUESTION NO: 4

Which actions will trigger partition rebalance for a consumer group? (select three)

A. Increase partitions of a topic

B. Remove a broker from the cluster

C. Add a new consumer to consumer group

D. A consumer in a consumer group shuts down
Add a broker to the cluster

ANSWER: A C D

Explanation:

:

Rebalance occurs when a new consumer is added, removed or consumer dies or paritions increased.

QUESTION NO: 5

To transform data from a Kafka topic to another one, I should use

A. Kafka Connect Sink

B. Kafka Connect Source

C. Consumer + Producer

D. Kafka Streams

ANSWER: D

Explanation:

:

Kafka Streams is a library for building streaming applications, specifically applications that transform input Kafka topics into output Kafka topics

QUESTION NO: 6

How do Kafka brokers ensure great performance between the producers and consumers?

(select two)

A. It compresses the messages as it writes to the disk

B. It leverages zero-copy optimisations to send data straight from the page-cache

C. It buffers the messages on disk, and sends messages from the disk reads

D. It transforms the messages into a binary format

E. It does not transform the messages

ANSWER: B E

Explanation:

:

Kafka transfers data with zero-copy and sends the raw bytes it receives from the producer straight to the consumer, leveraging the RAM available as page cache

QUESTION NO: 7

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

ANSWER: C

Explanation:

:

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

QUESTION NO: 8

A consumer wants to read messages from a specific partition of a topic. How can this be achieved?

A. Call subscribe(String topic, int partition) passing the topic and partition number as the arguments

B. Call assign() passing a Collection of TopicPartitions as the argument

C. Call subscribe() passing TopicPartition as the argument

ANSWER: B

Explanation:

:

assign() can be used for manual assignment of a partition to a consumer, in which case subscribe() must not be used. Assign() takes a collection of TopicPartition object as an argument
https://kafka.apache.org/23/javadoc/org/apache/kafka/clients/consumer/KafkaConsumer.ht ml#assign-java.util.Collection-

QUESTION NO: 9

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?

A. offset 45

B. offset 10

C. it will crash

D. offset 2311

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 NO: 10

What client protocol is supported for the schema registry? (select two)

A. HTTP

B. HTTPS

C. JDBC

D. Websocket

E. SASL

ANSWER: A B

Explanation:

:

clients can interact with the schema registry using the HTTP or HTTPS interface