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

How much should be the heap size of a broker in a production setup on a machine with 256 GB of RAM, in PLAINTEXT mode?

A. 4 GB

B. 128 GB

C. 16 GB

D. 512 MB

Correct Answer: A

In Kafka, a small heap size is needed, while the rest of the RAM goes automatically to the page cache (managed by the OS). The heap size goes slightly up if you need to enable SSL

QUESTION 2

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 3

You have a Zookeeper cluster that needs to be able to withstand the loss of 2 servers and still be able to function. What size should your Zookeeper cluster have?

A. 4

B. 5

C. 2

D. 3

E. 6

Correct Answer: B

Your Zookeeper cluster needs to have an odd number of servers, and must maintain a majority of servers up to be able to vote. Therefore, a 2N+1 zookeeper cluster can survive to N zookeeper being down, so here the right answer is N=2, 2*N+1=5

QUESTION 4

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 5

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 6

We want the average of all events in every five-minute window updated every minute. What kind of Kafka Streams window will be required on the stream?

A. Session window

B. Tumbling window

C. Sliding window

D. Hopping window

Correct Answer: D

A hopping window is defined by two propertiesthe window\\'s size and its advance interval (aka "hop"), e.g., a hopping window with a size 5 minutes and an advance interval of 1 minute.

QUESTION 7

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

A. HTTP

B. HTTPS

C. JDBC

D. Websocket

E. SASL

Correct Answer: AB

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

QUESTION 8

Which of these joins does not require input topics to be sharing the same number of partitions?

A. KStream-KTable join

B. KStream-KStream join

C. KStream-GlobalKTable

D. KTable-KTable join

Correct Answer: C

GlobalKTables have their datasets replicated on each Kafka Streams instance and therefore no repartitioning is required

QUESTION 9

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

QUESTION 10

You are using JDBC source connector to copy data from 2 tables to two 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. 6

B. 1

C. 2

D. 3

Correct Answer: C

we have two tables, so the max number of tasks is 2

QUESTION 11

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)

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

Correct Answer: BE

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 12

What exceptions may be caught by the following producer? (select two)

ProducerRecord record =

new ProducerRecord("topic1", "key1", "value1");

try {

producer.send(record);

} catch (Exception e) {

e.printStackTrace();

}

A. BrokerNotAvailableException

B. SerializationException

C. InvalidPartitionsException

D. BufferExhaustedException

Correct Answer: BD

These are the client side exceptions that may be encountered before message is sent to the broker, and before a future is returned by the .send() method.

QUESTION 13

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

A. 3

B. 2

C. 1

D. 6

Correct Answer: C

JDBC connector allows one task per table.

QUESTION 14

If I want to have an extremely high confidence that leaders and replicas have my data, I should use

A. acks=all, replication factor=2, min.insync.replicas=1

B. acks=1, replication factor=3, min.insync.replicas=2

C. acks=all, replication factor=3, min.insync.replicas=2

D. acks=all, replication factor=3, min.insync.replicas=1

Correct Answer: C

acks=all means the leader will wait for all in-sync replicas to acknowledge the record. Also the min in-sync replica setting specifies the minimum number of replicas that need to be in- sync for the partition to remain available for writes.

QUESTION 15

You want to send a message of size 3 MB to a topic with default message size configuration. How does KafkaProducer handle large messages?

A. KafkaProducer divides messages into sizes of max.request.size and sends them in order

B. KafkaProducer divides messages into sizes of message.max.bytes and sends them in order

C. MessageSizeTooLarge exception will be thrown, KafkaProducer will not retry and return exception immediately

D. MessageSizeTooLarge exception will be thrown, KafkaProducer retries until the number of retries are exhausted

Correct Answer: C

MessageSizeTooLarge is not a retryable exception.