



Free Questions for CCDAK by dumpshq

Shared by Lowe on 15-04-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

What is true about partitions? (select two)

Options:

- A- A broker can have a partition and its replica on its disk
- B- You cannot have more partitions than the number of brokers in your cluster
- C- A broker can have different partitions numbers for the same topic on its disk
- D- Only out of sync replicas are replicas, the remaining partitions that are in sync are also leader
- E- A partition has one replica that is a leader, while the other replicas are followers

Answer:

C, E

Explanation:

Only one of the replicas is elected as partition leader. And a broker can definitely hold many partitions from the same topic on its disk, try creating a topic with 12 partitions on one broker!

Question 2

Question Type: MultipleChoice

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

Options:

- 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 3

Question Type: MultipleChoice

There are 3 producers writing to a topic with 5 partitions. There are 10 consumers consuming from the topic as part of the same group. How many consumers will remain idle?

Options:

A- 10

B- 3

C- None

D- 5

Answer:

D

Explanation:

One consumer per partition assignment will keep 5 consumers idle.

Question 4

Question Type: MultipleChoice

Select the Kafka Streams joins that are always windowed joins.

Options:

- A- KStream-KStream join
- B- KTable-KTable join
- C- KStream-GlobalKTable
- D- KStream-KTable join

Answer:

A

Explanation:

See <https://docs.confluent.io/current/streams/developer-guide/dsl-api.html#joining>

Question 5

Question Type: MultipleChoice

```
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?

Options:

A- max.message.bytes=10000000

B- cleanup.policy=delete

C- compression.type=lz4

D- cleanup.policy=compact

Answer:

D

Explanation:

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 6

Question Type: MultipleChoice

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

Options:

- A- Future<ProducerRecord> object
- B- A Boolean indicating if the call succeeded
- C- Future<RecordMetadata> object
- D- Unit

Answer:

C

Explanation:

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

To Get Premium Files for CCDAK Visit

<https://www.p2pexams.com/products/ccdak>

For More Free Questions Visit

<https://www.p2pexams.com/confluent/pdf/ccdak>

