

DUMPS ARENA

Confluent Certified Developer for Apache Kafka Certification Examination

Confluent CCDAK

Version Demo

Total Demo Questions: 10

Total Premium Questions: 150

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

Which KSQL queries write to Kafka?

- A. COUNT and JOIN
- B. SHOW STREAMS and EXPLAIN statements
- C. CREATE STREAM WITH and CREATE TABLE WITH
- D. CREATE STREAM AS SELECT and CREATE TABLE AS SELECT

ANSWER: C D**Explanation:**

:

SHOW STREAMS and EXPLAIN statements run against the KSQL server that the

KSQL client is connected to. They don't communicate directly with Kafka. CREATE STREAM WITH and CREATE TABLE WITH write metadata to the KSQL command topic. Persistent queries based on CREATE STREAM AS SELECT and CREATE TABLE AS SELECT read and write to Kafka topics. Non-persistent queries based on SELECT that are stateless only read from Kafka topics, for example SELECT ,Ä¶ FROM foo WHERE ,Ä¶. Non-persistent queries that are stateful read and write to Kafka, for example, COUNT and JOIN. The data in Kafka is deleted automatically when you terminate the query with CTRL-C.

QUESTION NO: 2

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

ANSWER: A C**Explanation:**

:

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

QUESTION NO: 3

Which of the following Kafka Streams operators are stateless? (select all that apply)

- A. map
- B. filter
- C. flatmap
- D. branch
- E. groupBy
- F. aggregate

ANSWER: A B C D E**Explanation:**

:

See <https://kafka.apache.org/20/documentation/streams/developer-guide/dslapi.html#stateless-transformations>

QUESTION NO: 4

Which Kafka CLI should you use to consume from a topic?

- A. kafka-console-consumer
- B. kafka-topics
- C. kafka-console
- D. kafka-consumer-groups

ANSWER: A**Explanation:**

:

Example `kafka-console-consumer --bootstrap-server 127.0.0.1:9092 --topic test --frombeginning`

QUESTION NO: 5

Which of the following setting increases the chance of batching for a Kafka Producer?

- A. Increase batch.size
- B. Increase message.max.bytes
- C. Increase the number of producer threads
- D. Increase linger.ms

ANSWER: D

Explanation:

:

linger.ms forces the producer to wait to send messages, hence increasing the chance of creating batches

QUESTION NO: 6

I am producing Avro data on my Kafka cluster that is integrated with the Confluent Schema Registry. After a schema change that is incompatible, I know my data will be rejected.

Which component will reject the data?

- A. The Confluent Schema Registry
- B. The Kafka Broker
- C. The Kafka Producer itself
- D. Zookeeper

ANSWER: A

Explanation:

:

The Confluent Schema Registry is your safeguard against incompatible schema changes and will be the component that ensures no breaking schema evolution will be possible. Kafka Brokers do not look at your payload and your payload schema, and therefore will not reject data

QUESTION NO: 7

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 partitions increased.

QUESTION NO: 8

CORRECT TEXT

If I want to send binary data through the REST proxy to topic "test_binary", it needs to be base64 encoded. A consumer connecting directly into the Kafka topic A. "test_binary" will receive

- B. binary data
- C. avro data
- D. json data
- E. base64 encoded data, it will need to decode it

Answer: B

Explanation:

:

On the producer side, after receiving base64 data, the REST Proxy will convert it into bytes and then send that bytes payload to Kafka. Therefore consumers reading directly from Kafka will receive binary data.

- A. binary data
- B. avro data
- C. json data
- D. base64 encoded data, it will need to decode it

ANSWER: A

Explanation:

:

On the producer side, after receiving base64 data, the REST Proxy will convert it into bytes and then send that bytes payload to Kafka. Therefore consumers reading directly from Kafka will receive binary data.

QUESTION NO: 9

How often is log compaction evaluated?

- A. Every time a new partition is created
- B. Every time a segment is closed
- C. Every time a message is sent to Kafka
- D. Every time a message is flushed to disk

ANSWER: B**Explanation:**

:

Log compaction is evaluated every time a segment is closed. It will be triggered if enough data is "dirty" (see dirty ratio config)

QUESTION NO: 10

Which of the following Kafka Streams operators are stateful? (select all that apply)

- A. flatmap
- B. reduce
- C. joining
- D. count
- E. peek
- F. aggregate

ANSWER: B C D F**Explanation:**

:

See <https://kafka.apache.org/20/documentation/streams/developer-guide/dslapi.html#stateful-transformations>