

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.pass2lead.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

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

Correct Answer: CD

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 ,? FROM foo WHERE ,?. 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 2

How will you read all the messages from a topic in your KSQL query?

- A. KSQL reads from the beginning of a topic, by default.
- B. KSQL reads from the end of a topic. This cannot be changed.
- C. Use KSQL CLI to set auto.offset.reset property to earliest

Correct Answer: C

Consumers can set auto.offset.reset property to earliest to start consuming from beginning. For KSQL, SET `\\'auto.offset.reset\\'==\\'earliest\\'`;

QUESTION 3

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

Correct Answer: BCDF

See <https://kafka.apache.org/20/documentation/streams/developer-guide/dsl-api.html#statefultransformations>

QUESTION 4

Your producer is producing at a very high rate and the batches are completely full each time. How can you improve the producer throughput? (select two)

- A. Enable compression
- B. Disable compression
- C. Increase batch.size
- D. Decrease batch.size
- E. Decrease linger.ms
- F. Increase linger.ms

Correct Answer: AC

batch.size controls how many bytes of data to collect before sending messages to the Kafka broker. Set this as high as possible, without exceeding available memory. Enabling compression can also help make more compact batches and increase the throughput of your producer. Linger.ms will have no effect as the batches are already full

QUESTION 5

When using the Confluent Kafka Distribution, where does the schema registry reside?

- A. As a separate JVM component
- B. As an in-memory plugin on your Zookeeper cluster
- C. As an in-memory plugin on your Kafka Brokers
- D. As an in-memory plugin on your Kafka Connect Workers

Correct Answer: A

Schema registry is a separate application that provides RESTful interface for storing and retrieving Avro schemas.