

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.leads4pass.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

You are using JDBC source connector to copy data from 3 tables to three 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. 2
- B. 1
- C. 3
- D. 6

Correct Answer: A

here, we have three tables, but the max.tasks is 2, so that's the maximum number of tasks that will be created

QUESTION 2

There are two consumers C1 and C2 belonging to the same group G subscribed to topics T1 and T2. Each of the topics has 3 partitions. How will the partitions be assigned to consumers with Partition Assigner being Round Robin Assigner?

- A. C1 will be assigned partitions 0 and 2 from T1 and partition 1 from T2. C2 will have partition 1 from T1 and partitions 0 and 2 from T2.
- B. Two consumers cannot read from two topics at the same time
- C. C1 will be assigned partitions 0 and 1 from T1 and T2, C2 will be assigned partition 2 from T1 and T2.
- D. All consumers will read from all partitions

Correct Answer: A

The correct option is the only one where the two consumers share an equal number of partitions amongst the two topics of three partitions. An interesting article to read is <https://medium.com/@anyili0928/what-i-have-learned-from-kafka-partition-assignment-strategy-799fdf15d3ab>

QUESTION 3

What is the protocol used by Kafka clients to securely connect to the Confluent REST Proxy?

- A. Kerberos
- B. SASL
- C. HTTPS (SSL/TLS)
- D. HTTP

Correct Answer: C

TLS - but it is still called SSL.

QUESTION 4

```
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?

- A. max.message.bytes=10000000
- B. cleanup.policy=delete
- C. compression.type=lz4
- D. cleanup.policy=compact

Correct Answer: D

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 5

What is a generic unique id that I can use for messages I receive from a consumer?

- A. topic + partition + timestamp
- B. topic + partition + offset
- C. topic + timestamp

Correct Answer: B

(Topic,Partition,Offset) uniquely identifies a message in Kafka

[CCDAK Practice Test](#)

[CCDAK Study Guide](#)

[CCDAK Braindumps](#)