

2022 New Training Course CCDAK Tutorial Preparation Guide [Q34-Q48]



2022 New Training Course CCDAK Tutorial Preparation Guide Dumps of CCDAK Cover all the requirements of the Real Exam QUESTION 34

By default, which replica will be elected as a partition leader? (select two)

- * Preferred leader broker if it is in-sync and auto.leader.rebalance.enable=true
- * Any of the replicas
- * Preferred leader broker if it is in-sync and auto.leader.rebalance.enable=false
- * An in-sync replica

Preferred leader is a broker that was leader when topic was created. It is preferred because when partitions are first created, the leaders are balanced between brokers. Otherwise, any of the in-sync replicas (ISR) will be elected leader, as long as unclean.leader.election=false (by default)

QUESTION 35

```
StreamsBuilder builder = new StreamsBuilder();
```

```
KStream<String, String> textLines = builder.stream(word-count-input); KTable<String, Long> wordCounts =  
textLines
```

```
.mapValues(textLine -> textLine.toLowerCase())
```

```
.flatMapValues(textLine -> Arrays.asList(textLine.split("W+")))
```

```
.selectKey((key, word) -> word)
```

```
.groupByKey()
```

```
.count(Materialized.as(wordCounts));
```

```
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?

- * max.message.bytes=10000000
- * cleanup.policy=delete
- * compression.type=lz4
- * cleanup.policy=compact

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 36

Suppose you have 6 brokers and you decide to create a topic with 10 partitions and a replication factor of 3. The brokers 0 and 1 are on rack A, the brokers 2 and 3 are on rack B, and the brokers 4 and 5 are on rack C.

If the leader for partition 0 is on broker 4, and the first replica is on broker 2, which broker can host the last replica? (select two)

- * 6
- * 1
- * 2
- * 5
- * 0
- * 3

When you create a new topic, partitions replicas are spread across racks to maintain availability. Hence, the Rack A, which currently does not hold the topic partition, will be selected for the last replica

QUESTION 37

The exactly once guarantee in the Kafka Streams is for which flow of data?

- * Kafka => Kafka
- * Kafka => External
- * External => Kafka

Kafka Streams can only guarantee exactly once processing if you have a Kafka to Kafka topology.

QUESTION 38

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?

- * 6
- * 1
- * 2
- * 3

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

QUESTION 39

If you enable an SSL endpoint in Kafka, what feature of Kafka will be lost?

- * Cross-cluster mirroring
- * Support for Avro format
- * Zero copy
- * Exactly-once delivery

With SSL, messages will need to be encrypted and decrypted, by being first loaded into the JVM, so you lose the zero copy optimization. See more information here <https://twitter.com/ijuma/status/1161303431501324293?s=09>

QUESTION 40

A producer application in a developer machine was able to send messages to a Kafka topic. After copying the producer application into another developer's machine, the producer is able to connect to Kafka but unable to produce to the same Kafka topic because of an authorization issue. What is the likely issue?

- * Broker configuration needs to be changed to allow a different producer
- * You cannot copy a producer application from one machine to another
- * The Kafka ACL does not allow another machine IP
- * The Kafka Broker needs to be rebooted

ACLs take `Host` as a parameter, which represents an IP. It can be `*` (all IP), or a specific IP. Here, it's a specific IP as moving a producer to a different machine breaks the consumer, so the ACL needs to be updated

QUESTION 41

Select all that applies (select THREE)

- * `min.insync.replicas` is a producer setting
- * `acks` is a topic setting
- * `acks` is a producer setting
- * `min.insync.replicas` is a topic setting
- * `min.insync.replicas` matters regardless of the values of `acks`
- * `min.insync.replicas` only matters if `acks=all`

`acks` is a producer setting `min.insync.replicas` is a topic or broker setting and is only effective when `acks=all`

QUESTION 42

Where are the dynamic configurations for a topic stored?

- * In Zookeeper
- * In an internal Kafka topic `__topic_configs`
- * In `server.properties`
- * On the Kafka broker file system

Dynamic topic configurations are maintained in Zookeeper.

QUESTION 43

A topic has three replicas and you set `min.insync.replicas` to 2. If two out of three replicas are not available, what happens when a consume request is sent to broker?

- * Data will be returned from the remaining in-sync replica
- * An empty message will be returned
- * `NotEnoughReplicasException` will be returned
- * A new leader for the partition will be elected

With this configuration, a single in-sync replica is still readable, but not writeable if the producer using `acks=all`

QUESTION 44

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

- * `map`
- * `filter`
- * `flatMap`
- * `branch`
- * `groupBy`
- * `aggregate`

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

QUESTION 45

In Avro, removing a field that does not have a default is a ___ schema evolution

- * breaking
- * full
- * backward
- * forward

Clients with new schema will be able to read records saved with old schema.

QUESTION 46

A Zookeeper ensemble contains 5 servers. What is the maximum number of servers that can go missing and the ensemble still run?

- * 3
- * 4
- * 2
- * 1

majority consists of 3 zk nodes for 5 nodes zk cluster, so 2 can fail

QUESTION 47

A consumer has `auto.offset.reset=latest`, and the topic partition currently has data for offsets going from 45 to 2311. The consumer group never committed offsets for the topic before. Where will the consumer read from?

- * offset 2311
- * offset 0
- * offset 45
- * it will crash

Latest means that data retrievals will start from where the offsets currently end

QUESTION 48

Which of the following statements are true regarding the number of partitions of a topic?

- * The number of partitions in a topic cannot be altered
 - * We can add partitions in a topic by adding a broker to the cluster
 - * We can add partitions in a topic using the kafka-topics.sh command
 - * We can remove partitions in a topic by removing a broker
 - * We can remove partitions in a topic using the kafka-topics.sh command
- We can only add partitions to an existing topic, and it must be done using the kafka-topics.sh command

Sample Questions of CCDAK Dumps With 100% Exam Passing Guarantee:
<https://www.topexamcollection.com/CCDAK-vce-collection.html>