

# Confluent

## Exam Questions CCDAK

Confluent Certified Developer for Apache Kafka Certification Examination



#### NEW QUESTION 1

What is true about replicas ?

- A. Produce requests can be done to the replicas that are followers
- B. Produce and consume requests are load-balanced between Leader and Follower replicas
- C. Leader replica handles all produce and consume requests
- D. Follower replica handles all consume requests

**Answer:** C

**Explanation:**

Replicas are passive - they don't handle produce or consume request. Produce and consume requests get sent to the node hosting partition leader.

#### NEW QUESTION 2

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

- A. breaking
- B. full
- C. backward
- D. forward

**Answer:** C

**Explanation:**

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

#### NEW QUESTION 3

Where are the dynamic configurations for a topic stored?

- A. In Zookeeper
- B. In an internal Kafka topic topic\_configuratings
- C. In server.properties
- D. On the Kafka broker file system

**Answer:** A

**Explanation:**

Dynamic topic configurations are maintained in Zookeeper.

#### NEW QUESTION 4

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

**Answer:** A

**Explanation:**

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

#### NEW QUESTION 5

What Java library is KSQL based on?

- A. Kafka Streams
- B. REST Proxy
- C. Schema Registry
- D. Kafka Connect

**Answer:** A

**Explanation:**

KSQL is based on Kafka Streams and allows you to express transformations in the SQL language that get automatically converted to a Kafka Streams program in the backend

#### NEW QUESTION 6

Kafka is configured with following parameters - log.retention.hours = 168 log.retention.minutes = 168 log.retention.ms = 168 How long will the messages be retained for?

- A. Broker will not start due to bad configuration
- B. 168 ms
- C. 168 hours
- D. 168 minutes

**Answer:** B

**Explanation:**

If more than one similar config is specified, the smaller unit size will take precedence.

**NEW QUESTION 7**

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 --from-beginning`

**NEW QUESTION 8**

In Avro, removing or adding a field that has a default is a schema evolution

- A. full
- B. backward
- C. breaking
- D. forward

**Answer:** A

**Explanation:**

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

**NEW QUESTION 9**

You have a consumer group of 12 consumers and when a consumer gets killed by the process management system, rather abruptly, it does not trigger a graceful shutdown of your consumer. Therefore, it takes up to 10 seconds for a rebalance to happen. The business would like to have a 3 seconds rebalance time. What should you do? (select two)

- A. Increase `session.timeout.ms`
- B. Decrease `session.timeout.ms`
- C. Increase `heartbeat.interval.ms`
- D. decrease `max.poll.interval.ms`
- E. increase `max.poll.interval.ms`
- F. Decrease `heartbeat.interval.ms`

**Answer:** BE

**Explanation:**

`session.timeout.ms` must be decreased to 3 seconds to allow for a faster rebalance, and the heartbeat thread must be quicker, so we also need to decrease `heartbeat.interval.ms`

**NEW QUESTION 10**

Select all that applies (select THREE)

- A. `min.insync.replicas` is a producer setting
- B. `acks` is a topic setting
- C. `acks` is a producer setting
- D. `min.insync.replicas` is a topic setting
- E. `min.insync.replicas` matters regardless of the values of `acks`
- F. `min.insync.replicas` only matters if `acks=all`

**Answer:** CDF

**Explanation:**

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

**NEW QUESTION 10**

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:** ABCDE

**Explanation:**

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

#### NEW QUESTION 11

We would like to be in an at-most once consuming scenario. Which offset commit strategy would you recommend?

- A. Commit the offsets on disk, after processing the data
- B. Do not commit any offsets and read from beginning
- C. Commit the offsets in Kafka, after processing the data
- D. Commit the offsets in Kafka, before processing the data

**Answer:** D

#### Explanation:

Here, we must commit the offsets right after receiving a batch from a call to `.poll()`

#### NEW QUESTION 16

In the Kafka consumer metrics it is observed that fetch-rate is very high and each fetch is small. What steps will you take to increase throughput?

- A. Increase `fetch.max.wait`
- B. Increase `fetch.max.bytes`
- C. Decrease `fetch.max.bytes`
- D. Decrease `fetch.min.bytes`
- E. Increase `fetch.min.bytes`

**Answer:** E

#### Explanation:

This will allow consumers to wait and receive more bytes in each fetch request.

#### NEW QUESTION 20

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

- A. 3
- B. 4
- C. 2
- D. 1

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 24

```
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("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`

**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.

#### NEW QUESTION 27

An ecommerce website sells some custom made goods. What's the natural way of modeling this data in Kafka streams?

- A. Purchase as stream, Product as stream, Customer as stream
- B. Purchase as stream, Product as table, Customer as table
- C. Purchase as table, Product as table, Customer as table
- D. Purchase as stream, Product as table, Customer as stream

**Answer:** B

#### Explanation:

Mostly-static data is modeled as a table whereas business transactions should be modeled as a stream.

#### NEW QUESTION 29

Your streams application is reading from an input topic that has 5 partitions. You run 5 instances of your application, each with `num.streams.threads` set to 5. How many stream tasks will be created and how many will be active?

- A. 5 created, 1 active
- B. 5 created, 5 active
- C. 25 created, 25 active
- D. 25 created, 5 active

**Answer:** D

#### Explanation:

One partition is assigned a thread, so only 5 will be active, and 25 threads (i.e. tasks) will be created

#### NEW QUESTION 30

Which KSQL queries write to Kafka?

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

**Answer:** CD

#### Explanation:

SHOW STREAMS and EXPLAIN <query> statements run against the KSQL server that the KSQL client is connected to. They don't communicate directly with Kafka. CREATE STREAM WITH <topic> and CREATE TABLE WITH <topic> 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.

#### NEW QUESTION 33

You are running a Kafka Streams application in a Docker container managed by Kubernetes, and upon application restart, it takes a long time for the docker container to replicate the state and get back to processing the data. How can you improve dramatically the application restart?

- A. Mount a persistent volume for your RocksDB
- B. Increase the number of partitions in your inputs topic
- C. Reduce the Streams caching property
- D. Increase the number of Streams threads

**Answer:** A

#### Explanation:

Although any Kafka Streams application is stateless as the state is stored in Kafka, it can take a while and lots of resources to recover the state from Kafka. In order to speed up recovery, it is advised to store the Kafka Streams state on a persistent volume, so that only the missing part of the state needs to be recovered.

#### NEW QUESTION 35

What happens if you write the following code in your producer? `producer.send(producerRecord).get()`

- A. Compression will be increased
- B. Throughput will be decreased
- C. It will force all brokers in Kafka to acknowledge the producerRecord
- D. Batching will be increased

**Answer:** B

#### Explanation:

Using `Future.get()` to wait for a reply from Kafka will limit throughput.

#### NEW QUESTION 39

To get acknowledgement of writes to only the leader partition, we need to use the config...

- A. `acks=1`
- B. `acks=0`
- C. `acks=all`

**Answer:** A

#### Explanation:

Producers can set `acks=1` to get acknowledgement from partition leader only.

#### NEW QUESTION 42

If I want to send binary data through the REST proxy, it needs to be base64 encoded. Which component needs to encode the binary data into base 64?

- A. The Producer
- B. The Kafka Broker

- C. Zookeeper
- D. The REST Proxy

**Answer:** A

**Explanation:**

The REST Proxy requires to receive data over REST that is already base64 encoded, hence it is the responsibility of the producer

**NEW QUESTION 44**

To enhance compression, I can increase the chances of batching by using

- A. acks=all
- B. linger.ms=20
- C. batch.size=65536
- D. max.message.size=10MB

**Answer:** B

**Explanation:**

linger.ms forces the producer to wait before sending messages, hence increasing the chance of creating batches that can be heavily compressed.

**NEW QUESTION 45**

You have a Zookeeper cluster that needs to be able to withstand the loss of 2 servers and still be able to function. What size should your Zookeeper cluster have?

- A. 4
- B. 5
- C. 2
- D. 3
- E. 6

**Answer:** B

**Explanation:**

Your Zookeeper cluster needs to have an odd number of servers, and must maintain a majority of servers up to be able to vote. Therefore, a  $2N+1$  zookeeper cluster can survive to  $N$  zookeeper being down, so here the right answer is  $N=2$ ,  $2*N+1=5$

**NEW QUESTION 50**

What's a Kafka partition made of?

- A. One file and one index
- B. One file
- C. One file and two indexes per segment
- D. One file and two indexes

**Answer:** C

**Explanation:**

Kafka partitions are made of segments (usually each segment is 1GB), and each segment has two corresponding indexes (offset index and time index)

**NEW QUESTION 54**

How can you gracefully make a Kafka consumer to stop immediately polling data from Kafka and gracefully shut down a consumer application?

- A. Call `consumer.wakeup()` and catch a `WakeupException`
- B. Call `consumer.poll()` in another thread
- C. Kill the consumer thread

**Answer:** A

**Explanation:**

See <https://stackoverflow.com/a/37748336/3019499>

**NEW QUESTION 58**

A Zookeeper configuration has tickTime of 2000, initLimit of 20 and syncLimit of 5. What's the timeout value for followers to connect to Zookeeper?

- A. 20 sec
- B. 10 sec
- C. 2000 ms
- D. 40 sec

**Answer:** D

**Explanation:**

tick time is 2000 ms, and initLimit is the config taken into account when establishing a connection to Zookeeper, so the answer is  $2000 * 20 = 40000$  ms = 40s

**NEW QUESTION 63**

To import data from external databases, I should use

- A. Confluent REST Proxy
- B. Kafka Connect Sink
- C. Kafka Streams
- D. Kafka Connect Source

**Answer:** D

**Explanation:**

Kafka Connect Sink is used to export data from Kafka to external databases and Kafka Connect Source is used to import from external databases into Kafka.

**NEW QUESTION 66**

A consumer wants to read messages from partitions 0 and 1 of a topic topic1. Code snippet is shown below.

```
consumer.subscribe(Arrays.asList("topic1")); List<TopicPartition> pc = new ArrayList<>();  
pc.add(new PartitionTopic("topic1", 0));  
pc.add(new PartitionTopic("topic1", 1)); consumer.assign(pc);
```

- A. This works fine
- B. subscribe() will subscribe to the topic and assign() will assign partitions to the consumer.
- C. Throws IllegalStateException

**Answer:** B

**Explanation:**

subscribe() and assign() cannot be called by the same consumer, subscribe() is used to leverage the consumer group mechanism, while assign() is used to manually control partition assignment and reads assignment

**NEW QUESTION 69**

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?

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 70**

You are sending messages with keys to a topic. To increase throughput, you decide to increase the number of partitions of the topic. Select all that apply.

- A. All the existing records will get rebalanced among the partitions to balance load
- B. New records with the same key will get written to the partition where old records with that key were written
- C. New records may get written to a different partition
- D. Old records will stay in their partitions

**Answer:** CD

**Explanation:**

Increasing the number of partition causes new messages keys to get hashed differently, and breaks the guarantee "same keys goes to the same partition". Kafka logs are immutable and the previous messages are not re-shuffled

**NEW QUESTION 72**

A kafka topic has a replication factor of 3 and min.insync.replicas setting of 1. What is the maximum number of brokers that can be down so that a producer with acks=all can still produce to the topic?

- A. 3
- B. 2
- C. 1

**Answer:** C

**Explanation:**

Two brokers can go down, and one replica will still be able to receive and serve data

**NEW QUESTION 73**

In Avro, adding a field to a record without default is a schema evolution

- A. forward
- B. backward
- C. full
- D. breaking

**Answer:** A

**Explanation:**

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

**NEW QUESTION 77**

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?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 82**

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

- A. Kafka => Kafka
- B. Kafka => External
- C. External => Kafka

**Answer:** A

**Explanation:**

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

**NEW QUESTION 83**

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:** AC

**Explanation:**

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

**NEW QUESTION 86**

To allow consumers in a group to resume at the previously committed offset, I need to set the proper value for...

- A. `value.deserializer`
- B. `auto.offset.resets`
- C. `group.id`
- D. `enable.auto.commit`

**Answer:** C

**Explanation:**

Setting a `group.id` that's consistent across restarts will allow your consumers part of the same group to resume reading from where offsets were last committed for that group

**NEW QUESTION 91**

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### **CCDAK Practice Exam Features:**

- \* CCDAK Questions and Answers Updated Frequently
- \* CCDAK Practice Questions Verified by Expert Senior Certified Staff
- \* CCDAK Most Realistic Questions that Guarantee you a Pass on Your First Try
- \* CCDAK Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The CCDAK Practice Test Here](#)**