

Confluent

Exam Questions CCDAK

Confluent Certified Developer for Apache Kafka Certification Examination



NEW QUESTION 1

What data format isn't natively available with the Confluent REST Proxy?

- A. avro
- B. binary
- C. protobuf
- D. json

Answer: C

Explanation:

Protocol buffers isn't a natively supported type for the Confluent REST Proxy, but you may use the binary format instead

NEW QUESTION 2

A consumer application is using KafkaAvroDeserializer to deserialize Avro messages. What happens if message schema is not present in AvroDeserializer local cache?

- A. Throws SerializationException
- B. Fails silently
- C. Throws DeserializationException
- D. Fetches schema from Schema Registry

Answer: D

Explanation:

First local cache is checked for the message schema. In case of cache miss, schema is pulled from the schema registry. An exception will be thrown in the Schema Registry does not have the schema (which should never happen if you set it up properly)

NEW QUESTION 3

Where are the dynamic configurations for a topic stored?

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

Answer: A

Explanation:

Dynamic topic configurations are maintained in Zookeeper.

NEW QUESTION 4

A consumer starts and has auto.offset.reset=none, and the topic partition currently has data for offsets going from 45 to 2311. The consumer group has committed the offset 10 for the topic before. Where will the consumer read from?

- A. offset 45
- B. offset 10
- C. it will crash
- D. offset 2311

Answer: C

Explanation:

auto.offset.reset=none means that the consumer will crash if the offsets it's recovering from have been deleted from Kafka, which is the case here, as 10 < 45

NEW QUESTION 5

How will you set the retention for the topic named ,Ä??my-topic,Ä?? to 1 hour?

- A. Set the broker config log.retention.ms to 3600000
- B. Set the consumer config retention.ms to 3600000
- C. Set the topic config retention.ms to 3600000
- D. Set the producer config retention.ms to 3600000

Answer: C

Explanation:

retention.ms can be configured at topic level while creating topic or by altering topic. It shouldn't be set at the broker level (log.retention.ms) as this would impact all the topics in the cluster, not just the one we are interested in

NEW QUESTION 6

What exceptions may be caught by the following producer? (select two) `ProducerRecord<String, String> record = new ProducerRecord<>("topic1", "key1", "value1"); try { producer.send(record); } catch (Exception e) { e.printStackTrace(); }`

- A. BrokerNotAvailableException
- B. SerializationException
- C. InvalidPartitionsException
- D. BufferExhaustedException

Answer: BD

Explanation:

These are the client side exceptions that may be encountered before message is sent to the broker, and before a future is returned by the .send() method.

NEW QUESTION 7

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

- A. Cross-cluster mirroring
- B. Support for Avro format
- C. Zero copy
- D. Exactly-once delivery

Answer: C

Explanation:

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>

NEW QUESTION 8

When using plain JSON data with Connect, you see the following error message `org.apache.kafka.connect.errors.DataExceptionJsonDeserializer with schemas.enable requires "schema" and "payload" fields and may not contain additional fields`. How will you fix the error?

- A. Set `key.converter`, `value.converter` to `JsonConverter` and the schema registry url
- B. Use Single Message Transforms to add schema and payload fields in the message
- C. Set `key.converter.schemas.enable` and `value.converter.schemas.enable` to false
- D. Set `key.converter`, `value.converter` to `AvroConverter` and the schema registry url

Answer: C

Explanation:

You will need to set the `schemas.enable` parameters for the converter to false for plain text with no schema.

NEW QUESTION 9

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` Increase `linger.ms`

Answer: AC

Explanation:

`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

NEW QUESTION 10

A Zookeeper ensemble contains 3 servers. Over which ports the members of the ensemble should be able to communicate in default configuration? (select three)

- A. 2181
- B. 3888
- C. 443
- D. 2888
- E. 9092
- F. 80

Answer: ABD

Explanation:

2181 - client port, 2888 - peer port, 3888 - leader port

NEW QUESTION 10

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

Answer: B

Explanation:

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

NEW QUESTION 12

If I supply the setting `compression.type=snappy` to my producer, what will happen? (select two)

- A. The Kafka brokers have to de-compress the data
- B. The Kafka brokers have to compress the data
- C. The Consumers have to de-compress the data
- D. The Consumers have to compress the data
- E. The Producers have to compress the data

Answer: C

Explanation:

Kafka transfers data with zero copy and no transformation. Any transformation (including compression) is the responsibility of clients.

NEW QUESTION 13

We have a store selling shoes. What dataset is a great candidate to be modeled as a KTable in Kafka Streams?

- A. Money made until now
- B. The transaction stream
- C. Items returned
- D. Inventory contents right now

Answer: AC

Explanation:

Aggregations of stream are stored in table, whereas Streams must be modeled as a KStream to avoid data explosion

NEW QUESTION 17

A producer application was sending messages to a partition with a replication factor of 2 by connecting to Broker 1 that was hosting partition leader. If the Broker 1 goes down, what will happen?

- A. The producer will automatically produce to the broker that has been elected leader
- B. The topic will be unavailable
- C. The producer will stop working

Answer: A

Explanation:

Once the client connects to any broker, it is connected to the entire cluster and in case of leadership changes, the clients automatically do a Metadata Request to an available broker to find out who is the new leader for the topic. Hence the producer will automatically keep on producing to the correct Kafka Broker

NEW QUESTION 20

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 21

A kafka topic has a replication factor of 3 and `min.insync.replicas` setting of 2. How many brokers can go down before a producer with `acks=all` can't produce?

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

Answer: C

Explanation:

`acks=all` and `min.insync.replicas=2` means we must have at least 2 brokers up for the partition to be available

NEW QUESTION 26

In Java, Avro `SpecificRecords` classes are

- A. automatically generated from an Avro Schema
- B. written manually by the programmer
- C. automatically generated from an Avro Schema + a Maven / Gradle Plugin

Answer: C

Explanation:

SpecificRecord is created from generated record classes

NEW QUESTION 30

If a topic has a replication factor of 3...

- A. 3 replicas of the same data will live on 1 broker
- B. Each partition will live on 4 different brokers
- C. Each partition will live on 2 different brokers
- D. Each partition will live on 3 different brokers

Answer: D

Explanation:

Replicas are spread across available brokers, and each replica = one broker. RF 3 = 3 brokers

NEW QUESTION 33

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 36

To read data from a topic, the following configuration is needed for the consumers

- A. all brokers of the cluster, and the topic name
- B. any broker to connect to, and the topic name
- C. the list of brokers that have the data, the topic name and the partitions list
- D. any broker, and the list of topic partitions

Answer: B

Explanation:

All brokers can respond to Metadata request, so a client can connect to any broker in the cluster.

NEW QUESTION 38

There are 3 producers writing to a topic with 5 partitions. There are 10 consumers consuming from the topic as part of the same group. How many consumers will remain idle?

- A. 10
- B. 3
- C. None
- D. 5

Answer: D

Explanation:

One consumer per partition assignment will keep 5 consumers idle.

NEW QUESTION 39

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 43

How will you find out all the partitions without a leader?

- A. kafka-topics.sh --broker-list localhost:9092 --describe --under-replicated-partitions
- B. kafka-topics.sh --bootstrap-server localhost:2181 --describe --unavailable-partitions
- C. kafka-topics.sh --zookeeper localhost:2181 --describe --unavailable-partitions
- D. kafka-topics.sh --zookeeper localhost:2181 --describe --under-replicated-partitions

Answer: C

Explanation:

Please note that as of Kafka 2.2, the --zookeeper option is deprecated and you can now use kafka-topics.sh --bootstrap-server localhost:9092 --describe --unavailable-partitions

NEW QUESTION 48

A producer just sent a message to the leader broker for a topic partition. The producer used acks=1 and therefore the data has not yet been replicated to followers. Under which conditions will the consumer see the message?

- A. Right away
- B. When the message has been fully replicated to all replicas
- C. Never, the produce request will fail
- D. When the high watermark has advanced

Answer: D

Explanation:

The high watermark is an advanced Kafka concept, and is advanced once all the ISR replicates the latest offsets. A consumer can only read up to the value of the High Watermark (which can be less than the highest offset, in the case of acks=1)

NEW QUESTION 52

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.

NEW QUESTION 53

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
- E. Add a broker to the cluster

Answer: ACD

Explanation:

Rebalance occurs when a new consumer is added, removed or consumer dies or partitions increased.

NEW QUESTION 54

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 55

A topic receives all the orders for the products that are available on a commerce site. Two applications want to process all the messages independently - order fulfilment and monitoring. The topic has 4 partitions, how would you organise the consumers for optimal performance and resource usage?

- A. Create 8 consumers in the same group with 4 consumers for each application
- B. Create two consumers groups for two applications with 8 consumers in each
- C. Create two consumer groups for two applications with 4 consumers in each
- D. Create four consumers in the same group, one for each partition - two for fulfilment and two for monitoring

Answer: C

Explanation:

two partitions groups - one for each application so that all messages are delivered to both the application. 4 consumers in each as there are 4 partitions of the topic, and you cannot have more consumers per groups than the number of partitions (otherwise they will be inactive and wasting resources)

NEW QUESTION 57

We want the average of all events in every five-minute window updated every minute. What kind of Kafka Streams window will be required on the stream?

- A. Session window
- B. Tumbling window
- C. Sliding window
- D. Hopping window

Answer: D

Explanation:

A hopping window is defined by two properties the window's size and its advance interval (aka "hop"), e.g., a hopping window with a size 5 minutes and an advance interval of 1 minute.

NEW QUESTION 61

A consumer is configured with `enable.auto.commit=false`. What happens when `close()` is called on the consumer object?

- A. The uncommitted offsets are committed
- B. A rebalance in the consumer group will happen immediately
- C. The group coordinator will discover that the consumer stopped sending heartbeat
- D. It will cause rebalance after `session.timeout.ms`

Answer: B

Explanation:

Calling `close()` on consumer immediately triggers a partition rebalance as the consumer will not be available anymore.

NEW QUESTION 66

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 71

You are using JDBC source connector to copy data from a table to Kafka topic. There is one connector created with `max.tasks` equal to 2 deployed on a cluster of 3 workers. How many tasks are launched?

- A. 3
- B. 2
- C. 1
- D. 6

Answer: C

Explanation:

JDBC connector allows one task per table.

NEW QUESTION 73

.....

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 FirstTry
- * CCDAK Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

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