



Confluent

Exam Questions CCDAK

Confluent Certified Developer for Apache Kafka Certification Examination

NEW QUESTION 1

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 2

What client protocol is supported for the schema registry? (select two)

- A. HTTP
- B. HTTPS
- C. JDBC
- D. WebSocket
- E. SASL

Answer: AB

Explanation:

clients can interact with the schema registry using the HTTP or HTTPS interface

NEW QUESTION 3

In Kafka Streams, by what value are internal topics prefixed by?

- A. tasks-<number>
- B. application.id
- C. group.id
- D. kafka-streams-

Answer: B

Explanation:

In Kafka Streams, the application.id is also the underlying group.id for your consumers, and the prefix for all internal topics (repartition and state)

NEW QUESTION 4

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

Answer: A

Explanation:

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>

NEW QUESTION 5

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 6

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 7

A bank uses a Kafka cluster for credit card payments. What should be the value of the property `unclean.leader.election.enable`?

- A. FALSE
- B. TRUE

Answer: A

Explanation:

Setting `unclean.leader.election.enable` to true means we allow out-of-sync replicas to become leaders, we will lose messages when this occurs, effectively losing credit card payments and making our customers very angry.

NEW QUESTION 8

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 9

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 10

Two consumers share the same `group.id` (consumer group id). Each consumer will

- A. Read mutually exclusive offsets blocks on all the partitions
- B. Read all the data on mutual exclusive partitions
- C. Read all data from all partitions

Answer: B

Explanation:

Each consumer is assigned a different partition of the topic to consume.

NEW QUESTION 10

How will you find out all the partitions where one or more of the replicas for the partition are not in-sync with the leader?

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

Answer: D

NEW QUESTION 15

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 17

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 19

A consumer sends a request to commit offset 2000. There is a temporary communication problem, so the broker never gets the request and therefore never responds. Meanwhile, the consumer processed another batch and successfully committed offset 3000. What should you do?

- A. Add a new consumer to the group
- B. Use the kafka-consumer-group command to manually commit the offsets 2000 for the consumer group
- C. Restart the consumer
- D. Nothing

Answer: D

Explanation:

In this case, because the offset 3000 has been committed and all the messages between 0 and 3000 have all been processed, it is okay not to have committed offset 2000. The right answer is to do "nothing", this behaviour is acceptable

NEW QUESTION 24

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 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

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 32

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 37

Producing with a key allows to...

- A. Ensure per-record level security
- B. Influence partitioning of the producer messages
- C. Add more information to my message
- D. Allow a Kafka Consumer to subscribe to a (topic,key) pair and only receive that data

Answer: B

Explanation:

Keys are necessary if you require strong ordering or grouping for messages that share the same key. If you require that messages with the same key are always seen in the correct order, attaching a key to messages will ensure messages with the same key always go to the same partition in a topic. Kafka guarantees order within a partition, but not across partitions in a topic, so alternatively not providing a key - which will result in round-robin distribution across partitions - will not maintain such order.

NEW QUESTION 39

Which of the following setting increases the chance of batching for a Kafka Producer?

- A. Increase batch.size
- B. Increase message.max.bytes
- C. Increase the number of producer threads
- D. Increase linger.ms

Answer: D

Explanation:

linger.ms forces the producer to wait to send messages, hence increasing the chance of creating batches

NEW QUESTION 43

What is returned by a producer.send() call in the Java API?

- A. Future<ProducerRecord> object
- B. A Boolean indicating if the call succeeded
- C. Future<RecordMetadata> object
- D. Unit

Answer: C

Explanation:

See <https://kafka.apache.org/21/javadoc/org/apache/kafka/clients/producer/KafkaProducer.html>

NEW QUESTION 47

Partition leader election is done by

- A. The consumers
- B. The Kafka Broker that is the Controller
- C. Zookeeper
- D. Vote amongst the brokers

Answer: C

Explanation:

The Controller is a broker that is responsible for electing partition leaders

NEW QUESTION 50

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=1 can't produce?

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

Answer: D

Explanation:

min.insync.replicas does not impact producers when acks=1 (only when acks=all)

NEW QUESTION 51

How much should be the heap size of a broker in a production setup on a machine with 256 GB of RAM, in PLAINTEXT mode?

- A. 4 GB
- B. 128 GB
- C. 16 GB
- D. 512 MB

Answer: A

Explanation:

In Kafka, a small heap size is needed, while the rest of the RAM goes automatically to the page cache (managed by the OS). The heap size goes slightly up if you need to enable SSL

NEW QUESTION 52

There are 3 producers writing to a topic with 5 partitions. There are 5 consumers consuming from the topic. How many Controllers will be present in the cluster?

- A. 3
- B. 5
- C. 2
- D. 1

Answer: D

Explanation:

There is only one controller in a cluster at all times.

NEW QUESTION 53

Compaction is enabled for a topic in Kafka by setting log.cleanup.policy=compact. What is true about log compaction?

- A. After cleanup, only one message per key is retained with the first value
- B. Each message stored in the topic is compressed
- C. Kafka automatically de-duplicates incoming messages based on key hashes
- D. After cleanup, only one message per key is retained with the latest value Compaction changes the offset of messages

Answer: D

Explanation:

Log compaction retains at least the last known value for each record key for a single topic partition. All compacted log offsets remain valid, even if record at offset has been compacted away as a consumer will get the next highest offset.

NEW QUESTION 55

Which is an optional field in an Avro record?

- A. doc
- B. name
- C. namespace
- D. fields

Answer: A

Explanation:

doc represents optional description of message

NEW QUESTION 59

What is the default port that the KSQL server listens on?

- A. 9092
- B. 8088
- C. 8083
- D. 2181

Answer: B

Explanation:

Default port of KSQL server is 8088

NEW QUESTION 60

What isn't an internal Kafka Connect topic?

- A. connect-status
- B. connect-offsets
- C. connect-configs

D. connect-jars

Answer: D

Explanation:

connect-configs stores configurations, connect-status helps to elect leaders for connect, and connect-offsets store source offsets for source connectors

NEW QUESTION 61

What are the requirements for a Kafka broker to connect to a Zookeeper ensemble? (select two)

- A. Unique value for each broker's zookeeper.connect parameter
- B. Unique values for each broker's broker.id parameter
- C. All the brokers must share the same broker.id
- D. All the brokers must share the same zookeeper.connect parameter

Answer: BD

Explanation:

Each broker must have a unique broker id and connect to the same zk ensemble and root zNode

NEW QUESTION 64

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?

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

Answer: C

Explanation:

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

NEW QUESTION 67

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 70

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 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 75

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

Answer: BCDF

Explanation:

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

NEW QUESTION 80

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 81

What isn't a feature of the Confluent schema registry?

- A. Store avro data
- B. Enforce compatibility rules
- C. Store schemas

Answer: A

Explanation:

Data is stored on brokers.

NEW QUESTION 82

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 86

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 88

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 89

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

The rule "same key goes to the same partition" is true unless...

- A. the number of producer changes
- B. the number of kafka broker changes
- C. the number of partition changes
- D. the replication factor changes

Answer: C

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 93

Which of the following errors are retrievable from a producer perspective? (select two)

- A. MESSAGE_TOO_LARGE
- B. INVALID_REQUIRED_ACKS
- C. NOT_ENOUGH_REPLICAS
- D. NOT_LEADER_FOR_PARTITION
- E. TOPIC_AUTHORIZATION_FAILED

Answer: CD

Explanation:

Both of these are retrievable errors, others non-retrievable errors. See the full list of errors and their "retrievable" status here https://kafka.apache.org/protocol#protocol_error_codes

NEW QUESTION 98

Where are the ACLs stored in a Kafka cluster by default?

- A. Inside the broker's data directory
- B. Under Zookeeper node /kafka-acl/
- C. In Kafka topic kafka_acls
- D. Inside the Zookeeper's data directory

Answer: A

Explanation:

ACLs are stored in Zookeeper node /kafka-acls/ by default.

NEW QUESTION 101

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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](#)