

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

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 4

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 5

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?

- A. offset 2311
- B. offset 0
- C. offset 45
- D. it will crash

Answer: A

Explanation:

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

NEW QUESTION 6

Which of these joins does not require input topics to be sharing the same number of partitions?

- A. KStream-KTable join
- B. KStream-KStream join
- C. KStream-GlobalkTable
- D. KTable-KTable join

Answer: C

Explanation:

GlobalKTables have their datasets replicated on each Kafka Streams instance and therefore no repartitioning is required

NEW QUESTION 7

Where are KSQL-related data and metadata stored?

- A. Kafka Topics
- B. Zookeeper
- C. PostgreSQL database
- D. Schema Registry

Answer: A

Explanation:

metadata is stored in and built from the KSQL command topic. Each KSQL server has its own in-memory version of the metastore.

NEW QUESTION 8

You have a Kafka cluster and all the topics have a replication factor of 3. One intern at your company stopped a broker, and accidentally deleted all the data of that broker on the disk. What will happen if the broker is restarted?

- A. The broker will start, and other topics will also be deleted as the broker data on the disk got deleted
- B. The broker will start, and won't be online until all the data it needs to have is replicated from other leaders
- C. The broker will crash
- D. The broker will start, and won't have any dat
- E. If the broker comes leader, we have a data loss

Answer: B

Explanation:

Kafka replication mechanism makes it resilient to the scenarios where the broker lose data on disk, but can recover from replicating from other brokers. This makes Kafka amazing!

NEW QUESTION 9

A customer has many consumer applications that process messages from a Kafka topic. Each consumer application can only process 50 MB/s. Your customer wants to achieve a target throughput of 1 GB/s. What is the minimum number of partitions will you suggest to the customer for that particular topic?

- A. 10
- B. 20
- C. 1
- D. 50

Answer: B

Explanation:

each consumer can process only 50 MB/s, so we need at least 20 consumers consuming one partition so that $50 * 20 = 1000$ MB target is achieved.

NEW QUESTION 10

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 10

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 13

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 17

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 20

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 23

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 25

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 27

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 31

How do you create a topic named test with 3 partitions and 3 replicas using the Kafka CLI?

- A. bin/kafka-topics.sh --create --broker-list localhost:9092 --replication-factor 3 --partitions 3--topic test
- B. bin/kafka-topics-create.sh --zookeeper localhost:9092 --replication-factor 3 --partitions 3--topic test
- C. bin/kafka-topics.sh --create --bootstrap-server localhost:9092 --replication-factor 3 -- partitions 3 --topic test
- D. bin/kafka-topics.sh --create --bootstrap-server localhost:2181 --replication-factor 3 -- partitions 3 --topic test

Answer: C

Explanation:

As of Kafka 2.3, the kafka-topics.sh command can take --bootstrap-server localhost:9092 as an argument. You could also use the (now deprecated) option of --zookeeper localhost:2181.

NEW QUESTION 35

An ecommerce wesbite 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 39

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 42

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

Answer: C

Explanation:

TLS - but it is still called SSL.

NEW QUESTION 45

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 48

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 49

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 53

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 57

What is true about partitions? (select two)

- A. A broker can have a partition and its replica on its disk
- B. You cannot have more partitions than the number of brokers in your cluster
- C. A broker can have different partitions numbers for the same topic on its disk
- D. Only out of sync replicas are replicas, the remaining partitions that are in sync are also leader
- E. A partition has one replica that is a leader, while the other replicas are followers

Answer: CE

Explanation:

Only one of the replicas is elected as partition leader. And a broker can definitely hold many partitions from the same topic on its disk, try creating a topic with 12 partitions on one broker!

NEW QUESTION 61

A topic "sales" is being produced to in the Americas region. You are mirroring this topic using Mirror Maker to the European region. From there, you are only reading the topic for analytics purposes. What kind of mirroring is this?

- A. Passive-Passive
- B. Active-Active
- C. Active-Passive

Answer: C

Explanation:

This is active-passing as the replicated topic is used for read-only purposes only

NEW QUESTION 65

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 67

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 71

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 75

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 78

Which of the following event processing application is stateless? (select two)

- A. Read events from a stream and modifies them from JSON to Avro
- B. Publish the top 10 stocks each day
- C. Read log messages from a stream and writes ERROR events into a high-priority stream and the rest of the events into a low-priority stream
- D. Find the minimum and maximum stock prices for each day of trading

Answer: AC

Explanation:

Stateless means processing of each message depends only on the message, so converting from JSON to Avro or filtering a stream are both stateless operations

NEW QUESTION 81

Consumer failed to process record # 10 and succeeded in processing record # 11. Select the course of action that you should choose to guarantee at least once processing

- A. Commit offsets at 10
- B. Do not commit until successfully processing the record #10
- C. Commit offsets at 11

Answer: C

Explanation:

Here, you shouldn't commit offsets 11 or 10 as it would indicate that the message #10 has been processed successfully.

NEW QUESTION 82

How do Kafka brokers ensure great performance between the producers and consumers? (select two)

- A. It compresses the messages as it writes to the disk
- B. It leverages zero-copy optimisations to send data straight from the page-cache
- C. It buffers the messages on disk, and sends messages from the disk reads
- D. It transforms the messages into a binary format
- E. It does not transform the messages

Answer: BE

Explanation:

Kafka transfers data with zero-copy and sends the raw bytes it receives from the producer straight to the consumer, leveraging the RAM available as page cache

NEW QUESTION 83

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 87

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 90

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 95

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 96

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 101

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 106

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 110

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 113

Using the Confluent Schema Registry, where are Avro schema stored?

- A. In the Schema Registry embedded SQL database
- B. In the Zookeeper node `/schemas`
- C. In the message bytes themselves
- D. In the `_schemas` topic

Answer: D

Explanation:

The Schema Registry stores all the schemas in the `_schemas` Kafka topic

NEW QUESTION 118

You want to sink data from a Kafka topic to S3 using Kafka Connect. There are 10 brokers in the cluster, the topic has 2 partitions with replication factor of 3. How many tasks will you configure for the S3 connector?

- A. 10
- B. 6
- C. 3
- D. 2

Answer: D

Explanation:

You cannot have more sink tasks (= consumers) than the number of partitions, so 2.

NEW QUESTION 121

You want to send a message of size 3 MB to a topic with default message size configuration. How does KafkaProducer handle large messages?

- A. KafkaProducer divides messages into sizes of max.request.size and sends them in order
- B. KafkaProducer divides messages into sizes of message.max.bytes and sends them in order
- C. MessageSizeTooLarge exception will be thrown, KafkaProducer will not retry and return exception immediately
- D. MessageSizeTooLarge exception will be thrown, KafkaProducer retries until the number of retries are exhausted

Answer: C

Explanation:

MessageSizeTooLarge is not a retryable exception.

NEW QUESTION 122

How often is log compaction evaluated?

- A. Every time a new partition is created
- B. Every time a segment is closed
- C. Every time a message is sent to Kafka
- D. Every time a message is flushed to disk

Answer: B

Explanation:

Log compaction is evaluated every time a segment is closed. It will be triggered if enough data is "dirty" (see dirty ratio config)

NEW QUESTION 124

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 129

What's is true about Kafka brokers and clients from version 0.10.2 onwards?

- A. Clients and brokers must have the exact same version to be able to communicate
- B. A newer client can talk to a newer broker, but an older client cannot talk to a newer broker
- C. A newer client can talk to a newer broker, and an older client can talk to a newer broker
- D. A newer client can't talk to a newer broker, but an older client can talk to a newer broker

Answer: C

Explanation:

Kafka's new bidirectional client compatibility introduced in 0.10.2 allows this. Read more here<https://www.confluent.io/blog/upgrading-apache-kafka-clients-just-got-easier/>

NEW QUESTION 132

Which of the following errors are retrieable 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 retrieable errors, others non-retrieable errors. See the full list of errors and their "retrieable" status herehttps://kafka.apache.org/protocol#protocol_error_codes

NEW QUESTION 134

.....

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](#)