

Confluent

Exam Questions CCDAK

Confluent Certified Developer for Apache Kafka Certification Examination



NEW QUESTION 1

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)

- A. 6
- B. 1
- C. 2
- D. 5
- E. 3

Answer: BE

Explanation:

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

NEW QUESTION 2

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 3

You want to perform table lookups against a KTable everytime a new record is received from the KStream. What is the output of KStream-KTable join?

- A. KTable
- B. GlobalKTable
- C. You choose between KStream or KTable
- D. KStream

Answer: D

Explanation:

Here KStream is being processed to create another KStream.

NEW QUESTION 4

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 5

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 6

You are building a consumer application that processes events from a Kafka topic. What is the most important metric to monitor to ensure real-time processing?

- A. UnderReplicatedPartitions
- B. records-lag-max
- C. MessagesInPerSec
- D. BytesInPerSec

Answer: B

Explanation:

This metric shows the current lag (number of messages behind the broker)

NEW QUESTION 7

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

- A. Preferred leader broker if it is in-sync and `auto.leader.rebalance.enable=true`
- B. Any of the replicas
- C. Preferred leader broker if it is in-sync and `auto.leader.rebalance.enable=false`
- D. An in-sync replica

Answer: BD

Explanation:

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)

NEW QUESTION 8

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 9

A client connects to a broker in the cluster and sends a fetch request for a partition in a topic. It gets an exception Not Leader For Partition Exception in the response. How does client handle this situation?

- A. Get the Broker id from Zookeeper that is hosting the leader replica and send request to it
- B. Send metadata request to the same broker for the topic and select the broker hosting the leader replica
- C. Send metadata request to Zookeeper for the topic and select the broker hosting the leader replica
- D. Send fetch request to each Broker in the cluster

Answer: B

Explanation:

In case the consumer has the wrong leader of a partition, it will issue a metadata request. The Metadata request can be handled by any node, so clients know afterwards which broker are the designated leader for the topic partitions. Produce and consume requests can only be sent to the node hosting partition leader.

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:

`Examplekafka-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 11

You are doing complex calculations using a machine learning framework on records fetched from a Kafka topic. It takes more about 6 minutes to process a record

batch, and the consumer enters rebalances even though it's still running. How can you improve this scenario?

- A. Increase max.poll.interval.ms to 600000
- B. Increase heartbeat.interval.ms to 600000
- C. Increase session.timeout.ms to 600000
- D. Add consumers to the consumer group and kill them right away

Answer: A

Explanation:

Here, we need to change the setting max.poll.interval.ms (default 300000) to its double in order to tell Kafka a consumer should be considered dead if the consumer only if it hasn't called the .poll() method in 10 minutes instead of 5.

NEW QUESTION 13

Your topic is log compacted and you are sending a message with the key K and value null. What will happen?

- A. The broker will delete all messages with the key K upon cleanup
- B. The producer will throw a Runtime exception
- C. The broker will delete the message with the key K and null value only upon cleanup
- D. The message will get ignored by the Kafka broker

Answer: A

Explanation:

Sending a message with the null value is called a tombstone in Kafka and will ensure the log compacted topic does not contain any messages with the key K upon compaction

NEW QUESTION 16

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 18

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 19

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

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 27

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 28

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 33

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 37

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 40

To transform data from a Kafka topic to another one, I should use

- A. Kafka Connect Sink
- B. Kafka Connect Source
- C. Consumer + Producer
- D. Kafka Streams

Answer: D

Explanation:

Kafka Streams is a library for building streaming applications, specifically applications that transform input Kafka topics into output Kafka topics

NEW QUESTION 41

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 46

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 50

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 53

A Kafka producer application wants to send log messages to a topic that does not include any key. What are the properties that are mandatory to configure for the producer configuration? (select three)

- A. bootstrap.servers
- B. partition
- C. key.serializer
- D. value.serializer
- E. key
- F. value

Answer: ACD

Explanation:

Both key and value serializer are mandatory.

NEW QUESTION 56

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 58

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 62

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 66

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 70

A producer is sending messages with null key to a topic with 6 partitions using the `DefaultPartitioner`. Where will the messages be stored?

- A. Partition 5
- B. Any of the topic partitions
- C. The partition for the null key
- D. Partition 0

Answer: A

Explanation:

Message with no keys will be stored with round-robin strategy among partitions.

NEW QUESTION 73

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 78

Select the Kafka Streams joins that are always windowed joins.

- A. `KStream-KStream` join
- B. `KTable-KTable` join
- C. `KStream-GlobalKTable`
- D. `KStream-KTable` join

Answer: A

Explanation:

See <https://docs.confluent.io/current/streams/developer-guide/dsl-api.html#joining>

NEW QUESTION 82

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 86

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 87

What happens when broker.rack configuration is provided in broker configuration in Kafka cluster?

- A. You can use the same broker.id as long as they have different broker.rack configuration
- B. Replicas for a partition are placed in the same rack
- C. Replicas for a partition are spread across different racks
- D. Each rack contains all the topics and partitions, effectively making Kafka highly available

Answer: C

Explanation:

Partitions for newly created topics are assigned in a rack alternating manner, this is the only change broker.rack does

NEW QUESTION 90

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 95

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 97

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 99

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

- A. The number of partitions in a topic cannot be altered
- B. We can add partitions in a topic by adding a broker to the cluster
- C. We can add partitions in a topic using the kafka-topics.sh command
- D. We can remove partitions in a topic by removing a broker
- E. We can remove partitions in a topic using the kafka-topics.sh command

Answer: C

Explanation:

We can only add partitions to an existing topic, and it must be done using the kafka-topics.sh command

NEW QUESTION 102

What information isn't stored inside of Zookeeper? (select two)

- A. Schema Registry schemas
- B. Consumer offset
- C. ACL information
- D. Controller registration
- E. Broker registration info

Answer: B

Explanation:

Consumer offsets are stored in a Kafka topic consumer_offsets, and the Schema Registry stored schemas in the _schemas topic.

NEW QUESTION 106

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 111

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 115

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 116

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 121

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 123

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 125

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 127

A consumer starts and has `auto.offset.reset=latest`, and the topic partition currently has data for offsets going from 45 to 2311. The consumer group has committed the offset 643 for the topic before. Where will the consumer read from?

- A. it will crash
- B. offset 2311
- C. offset 643
- D. offset 45

Answer: C

Explanation:

The offsets are already committed for this consumer group and topic partition, so the property `auto.offset.reset` is ignored

NEW QUESTION 129

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 134

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 135

What is the disadvantage of request/response communication?

- A. Scalability
- B. Reliability
- C. Coupling
- D. Cost

Answer: C

Explanation:

Point-to-point (request-response) style will couple client to the server.

NEW QUESTION 137

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 139

What is not a valid authentication mechanism in Kafka?

- A. SASL/GSSAPI
- B. SASL/SCRAM
- C. SAML
- D. SSL

Answer: C

Explanation:

Learn more about security here <https://kafka.apache.org/documentation/#security>

NEW QUESTION 144

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 145

.....

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](#)