

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 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 3

To continuously export data from Kafka into a target database, I should use

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

Answer: C

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 4

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 5

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 6

If I want to have an extremely high confidence that leaders and replicas have my data, I should use

- A. acks=all, replication factor=2, min.insync.replicas=1

- B. acks=1, replication factor=3, min.insync.replicas=2
- C. acks=all, replication factor=3, min.insync.replicas=2
- D. acks=all, replication factor=3, min.insync.replicas=1

Answer: C

Explanation:

acks=all means the leader will wait for all in-sync replicas to acknowledge the record. Also the min in-sync replica setting specifies the minimum number of replicas that need to be in- sync for the partition to remain available for writes.

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

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 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

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 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 12

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 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 22

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

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 29

If I produce to a topic that does not exist, and the broker setting `auto.create.topic.enable=true`, what will happen?

- A. Kafka will automatically create the topic with 1 partition and 1 replication factor
- B. Kafka will automatically create the topic with the indicated producer settings `num.partitions` and `default.replication.factor`
- C. Kafka will automatically create the topic with the broker settings `num.partitions` and `default.replication.factor`
- D. Kafka will automatically create the topic with `num.partitions=#of brokers` and `replication.factor=3`

Answer: C

Explanation:

The broker settings comes into play when a topic is auto created

NEW QUESTION 32

Once sent to a topic, a message can be modified

- A. No

B. Yes

Answer: A

Explanation:

Kafka logs are append-only and the data is immutable

NEW QUESTION 33

Your manager would like to have topic availability over consistency. Which setting do you need to change in order to enable that?

- A. compression.type
- B. unclean.leader.election.enable
- C. min.insync.replicas

Answer: B

Explanation:

unclean.leader.election.enable=true allows non ISR replicas to become leader, ensuring availability but losing consistency as data loss will occur

NEW QUESTION 34

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 35

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 38

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 39

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 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

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 48

The kafka-console-consumer CLI, when used with the default options

- A. uses a random group id
- B. always uses the same group id
- C. does not use a group id

Answer: A

Explanation:

If a group is not specified, the kafka-console-consumer generates a random consumer group.

NEW QUESTION 49

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 51

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 56

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 58

There are five brokers in a cluster, a topic with 10 partitions and replication factor of 3, and a quota of producer_bytes_rate of 1 MB/sec has been specified for the client. What is the maximum throughput allowed for the client?

- A. 10 MB/s
- B. 0.33 MB/s
- C. 1 MB/s
- D. 5 MB/s

Answer: D

Explanation:

Each producer is allowed to produce @ 1MB/s to a broker. Max throughput 5 * 1MB, because we have 5 brokers.

NEW QUESTION 63

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 66

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 69

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 71

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 74

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 75

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 80

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 85

You are receiving orders from different customer in an "orders" topic with multiple partitions. Each message has the customer name as the key. There is a special customer named ABC that generates a lot of orders and you would like to reserve a partition exclusively for ABC. The rest of the message should be distributed among other partitions. How can this be achieved?

- A. Add metadata to the producer record
- B. Create a custom partitioner
- C. All messages with the same key will go the same partition, but the same partition may have messages with different key
- D. It is not possible to reserve
- E. Define a Kafka Broker routing rule

Answer: B

Explanation:

A Custom Partitioner allows you to easily customise how the partition number gets computed from a source message.

NEW QUESTION 87

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 91

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 93

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 95

An ecommerce website maintains two topics - a high volume "purchase" topic with 5 partitions and low volume "customer" topic with 3 partitions. You would like to do a stream- table join of these topics. How should you proceed?

- A. Repartition the purchase topic to have 3 partitions
- B. Repartition customer topic to have 5 partitions
- C. Model customer as a GlobalKTable
- D. Do a KStream / KTable join after a repartition step

Answer: C

Explanation:

In case of KStream-KStream join, both need to be co-partitioned. This restriction is not applicable in case of join with GlobalKTable, which is the most efficient here.

NEW QUESTION 97

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 101

The Controller is a broker that is... (select two)

- A. elected by Zookeeper ensemble
- B. is responsible for partition leader election
- C. elected by broker majority
- D. is responsible for consumer group rebalances

Answer: AB

Explanation:

Controller is a broker that in addition to usual broker functions is responsible for partition leader election. The election of that broker happens thanks to Zookeeper and at any time only one broker can be a controller

NEW QUESTION 105

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 106

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 108

To produce data to a topic, a producer must provide the Kafka client with...

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

Answer: D

Explanation:

All brokers can respond to a Metadata request, so a client can connect to any broker in the cluster and then figure out on its own which brokers to send data to.

NEW QUESTION 110

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 114

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 115

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 119

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 124

.....

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](#)