



Confluent

Exam Questions CCDAK

Confluent Certified Developer for Apache Kafka Certification Examination

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

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

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

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

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

Answer: BD

Explanation:

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

NEW QUESTION 4

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 5

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 6

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 7

There are 3 brokers in the cluster. You want to create a topic with a single partition that is resilient to one broker failure and one broker maintenance. What is the replication factor will you specify while creating the topic?

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

Answer: B

Explanation:

1 is not possible as it doesn't provide resilience to failure, 2 is not enough as if we take a broker down for maintenance, we cannot tolerate a broker failure, and 6 is impossible as we only have 3 brokers (RF cannot be greater than the number of brokers). Here the correct answer is 3

NEW QUESTION 8

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 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

```
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 19

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 23

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 26

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 31

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 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 36

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 41

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 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 47

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 51

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 56

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 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 59

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 63

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 67

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 69

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 70

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 71

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 75

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 80

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

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

Answer: C

Explanation:

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

NEW QUESTION 83

What is the risk of increasing `max.in.flight.requests.per.connection` while also enabling retries in a producer?

- A. At least once delivery is not guaranteed
- B. Message order not preserved
- C. Reduce throughput
- D. Less resilient

Answer: B

Explanation:

Some messages may require multiple retries. If there are more than 1 requests in flight, it may result in messages received out of order. Note an exception to this rule is if you enable the producer setting `enable.idempotence=true` which takes care of the out of ordering case on its own.

See <https://issues.apache.org/jira/browse/KAFKA-5494>

NEW QUESTION 88

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 93

A producer application in a developer machine was able to send messages to a Kafka topic. After copying the producer application into another developer's machine, the producer is able to connect to Kafka but unable to produce to the same Kafka topic because of an authorization issue. What is the likely issue?

- A. Broker configuration needs to be changed to allow a different producer
- B. You cannot copy a producer application from one machine to another
- C. The Kafka ACL does not allow another machine IP
- D. The Kafka Broker needs to be rebooted

Answer: C

Explanation:

ACLs take "Host" as a parameter, which represents an IP. It can be * (all IP), or a specific IP. Here, it's a specific IP as moving a producer to a different machine breaks the consumer, so the ACL needs to be updated

NEW QUESTION 95

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 99

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

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 104

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 106

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 111

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 112

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 114

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 115

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 120

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 122

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 127

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 131

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 134

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 139

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 143

.....

Relate Links

100% Pass Your CCDAK Exam with ExamBible Prep Materials

<https://www.exambible.com/CCDAK-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>