



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

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

is KSQL ANSI SQL compliant?

- A. Yes
- B. No

Answer: B

Explanation:

KSQL is not ANSI SQL compliant, for now there are no defined standards on streaming SQL languages

NEW QUESTION 3

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 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

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 7

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

- A. FALSE
- B. TRUE

Answer: A

Explanation:

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

NEW QUESTION 8

When using the Confluent Kafka Distribution, where does the schema registry reside?

- A. As a separate JVM component
- B. As an in-memory plugin on your Zookeeper cluster
- C. As an in-memory plugin on your Kafka Brokers
- D. As an in-memory plugin on your Kafka Connect Workers

Answer: A

Explanation:

Schema registry is a separate application that provides RESTful interface for storing and retrieving Avro schemas.

NEW QUESTION 9

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 10

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 14

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 19

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 24

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 26

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 30

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

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 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 42

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

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 51

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 54

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 57

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 61

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

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

Answer: D

Explanation:

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

NEW QUESTION 62

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 66

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 70

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 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 72

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 73

A Zookeeper configuration has tickTime of 2000, initLimit of 20 and syncLimit of 5. What's the timeout value for followers to connect to Zookeeper?

- A. 20 sec
- B. 10 sec
- C. 2000 ms
- D. 40 sec

Answer: D

Explanation:

tick time is 2000 ms, and initLimit is the config taken into account when establishing a connection to Zookeeper, so the answer is $2000 * 20 = 40000$ ms = 40s

NEW QUESTION 77

To import data from external databases, I should use

- A. Confluent REST Proxy
- B. Kafka Connect Sink
- C. Kafka Streams
- D. Kafka Connect Source

Answer: D

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 80

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 85

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 86

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 87

You are using JDBC source connector to copy data from 2 tables to two Kafka topics. There is one connector created with max.tasks equal to 2 deployed on a cluster of 3 workers. How many tasks are launched?

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

Answer: C

Explanation:

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

NEW QUESTION 91

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 93

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 95

A kafka topic has a replication factor of 3 and min.insync.replicas setting of 1. What is the maximum number of brokers that can be down so that a producer with acks=all can still produce to the topic?

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

Answer: C

Explanation:

Two brokers can go down, and one replica will still be able to receive and serve data

NEW QUESTION 97

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 99

A topic has three replicas and you set `min.insync.replicas` to 2. If two out of three replicas are not available, what happens when a consume request is sent to broker?

- A. Data will be returned from the remaining in-sync replica
- B. An empty message will be returned
- C. `NotEnoughReplicasException` will be returned
- D. A new leader for the partition will be elected

Answer: A

Explanation:

With this configuration, a single in-sync replica is still readable, but not writeable if the producer using `acks=all`

NEW QUESTION 101

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 105

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 106

Which of the following is true regarding thread safety in the Java Kafka Clients?

- A. One Producer can be safely used in multiple threads
- B. One Consumer can be safely used in multiple threads
- C. One Consumer needs to run in one thread
- D. One Producer needs to be run in one thread

Answer: AC

Explanation:

`KafkaConsumer` is not thread-safe, `KafkaProducer` is thread safe.

NEW QUESTION 107

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 108

Select all the way for one consumer to subscribe simultaneously to the following topics - `topic.history`, `topic.sports`, `topic.politics`? (select two)

- A. `consumer.subscribe(Pattern.compile("topic\\..*"));`
- B. `consumer.subscribe("topic.history"); consumer.subscribe("topic.sports"); consumer.subscribe("topic.politics");`
- C. `consumer.subscribePrefix("topic.");`
- D. `consumer.subscribe(Arrays.asList("topic.history", "topic.sports", "topic.politics"));`

Answer: AD

Explanation:

Multiple topics can be passed as a list or regex pattern.

NEW QUESTION 113

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 114

.....

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/>