



Free Questions for [CCDAK](#) by [go4braindumps](#)

Shared by [Hopkins](#) on [29-01-2024](#)

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)

Question 1

Question Type: MultipleChoice

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

Options:

- 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.

Question 2

Question Type: MultipleChoice

What is not a valid authentication mechanism in Kafka?

Options:

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

Answer:

C

Explanation:

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

Question 3

Question Type: MultipleChoice

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?

Options:

- 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)

Question 4

Question Type: MultipleChoice

When `auto.create.topics.enable` is set to true in Kafka configuration, what are the circumstances under which a Kafka broker automatically creates a topic? (select three)

Options:

- A- Client requests metadata for a topic
- B- Consumer reads message from a topic
- C- Client alters number of partitions of a topic
- D- Producer sends message to a topic

Answer:

A, B, D

Explanation:

A kafka broker automatically creates a topic under the following circumstances- When a producer starts writing messages to the topic -
When a consumer starts reading messages from the topic - When any client requests metadata for the topic

Question 5

Question Type: MultipleChoice

When is the onCompletion() method called?

```
private class ProducerCallback implements Callback {  
  
    @Override  
  
    public void onCompletion(RecordMetadata recordMetadata, Exception e) {  
  
        if (e != null) {  
  
            e.printStackTrace();  
  
        }  
  
    }  
  
}
```

```
}  
  
ProducerRecord record =  
  
new ProducerRecord("topic1", "key1", "value1");  
  
producer.send(record, new ProducerCallback());
```

Options:

- A- When the message is partitioned and batched successfully
- B- When message is serialized successfully
- C- When the broker response is received
- D- When send() method is called

Answer:

C

Explanation:

Callback is invoked when a broker response is received.

Question 6

Question Type: MultipleChoice

Compaction is enabled for a topic in Kafka by setting `log.cleanup.policy=compact`. What is true about log compaction?

Options:

- 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.

Question 7

Question Type: MultipleChoice

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

Options:

- 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.

Question 8

Question Type: MultipleChoice

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?

Options:

- 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.

Question 9

Question Type: MultipleChoice

A kafka topic has a replication factor of 3 and min.insync.replicas setting of 2. How many brokers can go down before a producer with acks=all can't produce?

Options:

A- 0

B- 2

C- 1

D- 3

Answer:

C

Explanation:

acks=all and min.insync.replicas=2 means we must have at least 2 brokers up for the partition to be available

Question 10

Question Type: MultipleChoice

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?

Options:

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 $N=2$, $2*N+1=5$

To Get Premium Files for CCDAK Visit

<https://www.p2pexams.com/products/ccdak>

For More Free Questions Visit

<https://www.p2pexams.com/confluent/pdf/ccdak>

