

## Exam Questions CCDAK

Confluent Certified Developer for Apache Kafka Certification Examination

<https://www.2passeasy.com/dumps/CCDAK/>



#### 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

In Avro, removing a field that does not have a default is a schema evolution

- A. breaking
- B. full
- C. backward
- D. forward

**Answer:** C

#### Explanation:

Clients with new schema will be able to read records saved with old schema.

#### NEW QUESTION 4

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 5

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 6

There are two consumers C1 and C2 belonging to the same group G subscribed to topics T1 and T2. Each of the topics has 3 partitions. How will the partitions be assigned to consumers with Partition Assigner being Round Robin Assigner?

- A. C1 will be assigned partitions 0 and 2 from T1 and partition 1 from T2. C2 will have partition 1 from T1 and partitions 0 and 2 from T2.

- B. Two consumers cannot read from two topics at the same time
- C. C1 will be assigned partitions 0 and 1 from T1 and T2, C2 will be assigned partition 2 from T1 and T2.
- D. All consumers will read from all partitions

**Answer:** A

**Explanation:**

The correct option is the only one where the two consumers share an equal number of partitions amongst the two topics of three partitions. An interesting article to read is <https://medium.com/@anyili0928/what-i-have-learned-from-kafka-partition-assignment-strategy-799fdf15d3ab>

**NEW QUESTION 7**

A consumer has `auto.offset.reset=latest`, and the topic partition currently has data for offsets going from 45 to 2311. The consumer group never committed offsets for the topic before. Where will the consumer read from?

- A. offset 2311
- B. offset 0
- C. offset 45
- D. it will crash

**Answer:** A

**Explanation:**

Latest means that data retrievals will start from where the offsets currently end

**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**

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 10**

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 15

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 19

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 20

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 22

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 25

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 30

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 33

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 36

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 40

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 44

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 45

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?

- A. 2
- B. 1
- C. 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

#### NEW QUESTION 48

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 53

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 55

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 58

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 61

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 66

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 67

How does a consumer commit offsets in Kafka?

- A. It directly sends a message to the consumer\_offsets topic
- B. It interacts with the Group Coordinator broker
- C. It directly commits the offsets in Zookeeper

**Answer:** B

#### Explanation:

Consumers do not directly write to the consumer\_offsets topic, they instead interact with a broker that has been elected to manage that topic, which is the Group Coordinator broker

#### NEW QUESTION 72

To enhance compression, I can increase the chances of batching by using

- A. acks=all
- B. linger.ms=20
- C. batch.size=65536
- D. max.message.size=10MB

**Answer:** B

#### Explanation:

linger.ms forces the producer to wait before sending messages, hence increasing the chance of creating batches that can be heavily compressed.

#### NEW QUESTION 77

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 79

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 82

You are using JDBC source connector to copy data from 3 tables to three 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. 2
- B. 1
- C. 3
- D. 6

**Answer:** A

#### Explanation:

here, we have three tables, but the max.tasks is 2, so that's the maximum number of tasks that will be created

#### NEW QUESTION 86

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 90**

How will you read all the messages from a topic in your KSQL query?

- A. KSQL reads from the beginning of a topic, by default.
- B. KSQL reads from the end of a topic
- C. This cannot be changed.
- D. Use KSQL CLI to set `auto.offset.reset` property to earliest

**Answer:** C

**Explanation:**

Consumers can set `auto.offset.reset` property to earliest to start consuming from beginning. For KSQL, SET `'auto.offset.reset'='earliest'`;

**NEW QUESTION 92**

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 95**

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 97**

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 98**

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 99**

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 101**

I am producing Avro data on my Kafka cluster that is integrated with the Confluent Schema Registry. After a schema change that is incompatible, I know my data will be rejected. Which component will reject the data?

- A. The Confluent Schema Registry
- B. The Kafka Broker
- C. The Kafka Producer itself
- D. Zookeeper

**Answer:** A

**Explanation:**

The Confluent Schema Registry is your safeguard against incompatible schema changes and will be the component that ensures no breaking schema evolution will be possible. Kafka Brokers do not look at your payload and your payload schema, and therefore will not reject data

**NEW QUESTION 102**

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)

- 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:** ABD

**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

**NEW QUESTION 103**

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 104**

In Kafka, every broker... (select three)

- A. contains all the topics and all the partitions
- B. knows all the metadata for all topics and partitions
- C. is a controller
- D. knows the metadata for the topics and partitions it has on its disk
- E. is a bootstrap broker
- F. contains only a subset of the topics and the partitions

**Answer:** BEF

**Explanation:**

Kafka topics are divided into partitions and spread across brokers. Each brokers knows about all the metadata and each broker is a bootstrap broker, but only one of them is elected controller

**NEW QUESTION 106**

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 109**

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 111**

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 116**

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 121**

When is the onCompletion() method called?

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

- A. `e.printStackTrace();}}` `ProducerRecord<String, String> record = new ProducerRecord<>("topic1", "key1", "value1"); producer.send(record, new ProducerCallback());`
- B. When the message is partitioned and batched successfully
- C. When message is serialized successfully
- D. When the broker response is received
- E. When send() method is called

**Answer:** C

**Explanation:**

Callback is invoked when a broker response is received.

**NEW QUESTION 122**

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 124**

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 125**

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 128**

Which of the following Kafka Streams operators are stateful? (select all that apply)

- A. flatmap
- B. reduce
- C. joining
- D. count
- E. peek
- F. aggregate

**Answer:** BCDF

**Explanation:**

See <https://kafka.apache.org/20/documentation/streams/developer-guide/dsl-api.html#stateful-transformations>

**NEW QUESTION 132**

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 133**

We want the average of all events in every five-minute window updated every minute. What kind of Kafka Streams window will be required on the stream?

- A. Session window
- B. Tumbling window
- C. Sliding window
- D. Hopping window

**Answer:** D

**Explanation:**

A hopping window is defined by two properties the window's size and its advance interval (aka "hop"), e.g., a hopping window with a size 5 minutes and an advance interval of 1 minute.

**NEW QUESTION 138**

In Avro, adding an element to an enum without a default is a schema evolution

- A. breaking
- B. full
- C. backward
- D. forward

**Answer:** A

**Explanation:**

Since Confluent 5.4.0, Avro 1.9.1 is used. Since default value was added to enum complex type, the schema resolution changed from: (<1.9.1) if both are enums:\*\* if the writer's symbol is not present in the reader's enum, then an error is signalled. \*\*(>=1.9.1) if both are enums: if the writer's symbol is not present in the reader's enum and the reader has a default value, then that value is used, otherwise an error is signalled.

**NEW QUESTION 141**

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 145**

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 149**

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 154**

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 155**

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 157**

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual CCDAK Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the CCDAK Product From:

<https://www.2passeasy.com/dumps/CCDAK/>

### Money Back Guarantee

#### 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 FirstTry
- \* CCDAK Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year