

## 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

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

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 5

A consumer application is using KafkaAvroDeserializer to deserialize Avro messages. What happens if message schema is not present in AvroDeserializer local cache?

- A. Throws SerializationException
- B. Fails silently
- C. Throws DeserializationException
- D. Fetches schema from Schema Registry

**Answer:** D

#### Explanation:

First local cache is checked for the message schema. In case of cache miss, schema is pulled from the schema registry. An exception will be thrown in the Schema Registry does not have the schema (which should never happen if you set it up properly)

#### 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**

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**

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?

- 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)

**NEW QUESTION 9**

A customer has many consumer applications that process messages from a Kafka topic. Each consumer application can only process 50 MB/s. Your customer wants to achieve a target throughput of 1 GB/s. What is the minimum number of partitions will you suggest to the customer for that particular topic?

- A. 10
- B. 20
- C. 1
- D. 50

**Answer:** B

**Explanation:**

each consumer can process only 50 MB/s, so we need at least 20 consumers consuming one partition so that  $50 * 20 = 1000$  MB target is achieved.

**NEW QUESTION 10**

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 10**

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:**

Examplekafka-console-consumer --bootstrap-server 127.0.0.1:9092 --topic test --from- beginning

### NEW QUESTION 13

Which of the following is not an Avro primitive type?

- A. string
- B. long
- C. int
- D. date
- E. null

**Answer:** D

#### Explanation:

date is a logical type

### NEW QUESTION 15

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 20

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

- A. map
- B. filter
- C. flatmap
- D. branch
- E. groupBy
- F. aggregate

**Answer:** ABCDE

#### Explanation:

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

### NEW QUESTION 21

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 23

```
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 28

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 29

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 34

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 35

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 40

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 41

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 45**

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 47**

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 52**

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 55**

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 58**

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 61

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 66

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 67

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 68

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 72

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 75

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 76**

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 79**

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 84**

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 86**

CORRECT TEXT

If I want to send binary data through the REST proxy to topic "test\_binary", it needs to be base64 encoded. A consumer connecting directly into the Kafka topic

A. "test\_binary" will receive

B. binary data

C. avro data

D. json data

E. base64 encoded data, it will need to decode it

**Answer:** B

**Explanation:**

On the producer side, after receiving base64 data, the REST Proxy will convert it into bytes and then send that bytes payload to Kafka. Therefore consumers reading directly from Kafka will receive binary data.

**NEW QUESTION 90**

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 93**

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 98**

What isn't a feature of the Confluent schema registry?

- A. Store avro data
- B. Enforce compatibility rules
- C. Store schemas

**Answer:** A

**Explanation:**

Data is stored on brokers.

**NEW QUESTION 102**

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 106**

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 107**

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 110**

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 112**

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](https://kafka.apache.org/protocol#protocol_error_codes)

**NEW QUESTION 114**

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 118**

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 120**

.....

## 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