



Confluent

Exam Questions CCDAK

Confluent Certified Developer for Apache Kafka Certification Examination

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

```
while (true) {  
    ConsumerRecords<String, String> records = consumer.poll(100); try {  
        consumer.commitSync();  
    } catch (CommitFailedException e) { log.error("commit failed", e)  
    }  
    for (ConsumerRecord<String, String> record records)  
    {  
        System.out.printf("topic = %s, partition = %s, offset =  
        %d, customer = %s, country = %s ",  
        record.topic(), record.partition(), record.offset(), record.key(), record.value());  
    }  
}
```

What kind of delivery guarantee this consumer offers?

- A. Exactly-once
- B. At-least-once
- C. At-most-once

Answer: C

Explanation:

Here offset is committed before processing the message. If consumer crashes before processing the message, message will be lost when it comes back up.

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

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 6

Which of these joins does not require input topics to be sharing the same number of partitions?

- A. KStream-KTable join
- B. KStream-KStream join
- C. KStream-GlobalKTable
- D. KTable-KTable join

Answer: C

Explanation:

GlobalKTables have their datasets replicated on each Kafka Streams instance and therefore no repartitioning is required

NEW QUESTION 7

Where are KSQL-related data and metadata stored?

- A. Kafka Topics
- B. Zookeeper
- C. PostgreSQL database
- D. Schema Registry

Answer: A

Explanation:

metadata is stored in and built from the KSQL command topic. Each KSQL server has its own in-memory version of the metastore.

NEW QUESTION 8

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 9

You are doing complex calculations using a machine learning framework on records fetched from a Kafka topic. It takes more about 6 minutes to process a record batch, and the consumer enters rebalances even though it's still running. How can you improve this scenario?

- A. Increase max.poll.interval.ms to 600000
- B. Increase heartbeat.interval.ms to 600000
- C. Increase session.timeout.ms to 600000
- D. Add consumers to the consumer group and kill them right away

Answer: A

Explanation:

Here, we need to change the setting max.poll.interval.ms (default 300000) to its double in order to tell Kafka a consumer should be considered dead if the consumer only if it hasn't called the .poll() method in 10 minutes instead of 5.

NEW QUESTION 10

We would like to be in an at-most once consuming scenario. Which offset commit strategy would you recommend?

- A. Commit the offsets on disk, after processing the data
- B. Do not commit any offsets and read from beginning
- C. Commit the offsets in Kafka, after processing the data
- D. Commit the offsets in Kafka, before processing the data

Answer: D

Explanation:

Here, we must commit the offsets right after receiving a batch from a call to .poll()

NEW QUESTION 10

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 12

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 15

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 17

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 18

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 22

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 26

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 29

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 30

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 31

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 36

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 40

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 43

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 48

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 52

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 55

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 58

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 $2000 * 20 = 40000 \text{ ms} = 40\text{s}$

NEW QUESTION 59

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 64

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 65

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 67

A consumer wants to read messages from partitions 0 and 1 of a topic topic1. Code snippet is shown below.

```
consumer.subscribe(Arrays.asList("topic1")); List<TopicPartition> pc = new ArrayList<>();  
pc.add(new PartitionTopic("topic1", 0));  
pc.add(new PartitionTopic("topic1", 1)); consumer.assign(pc);
```

- A. This works fin
- B. subscribe() will subscribe to the topic and assign() will assign partitions to the consumer.
- C. Throws IllegalStateException

Answer: B

Explanation:

subscribe() and assign() cannot be called by the same consumer, subscribe() is used to leverage the consumer group mechanism, while assign() is used to manually control partition assignment and reads assignment

NEW QUESTION 72

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 73

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 74

What isn't an internal Kafka Connect topic?

- A. connect-status
- B. connect-offsets
- C. connect-configs
- D. connect-jars

Answer: D

Explanation:

connect-configs stores configurations, connect-status helps to elect leaders for connect, and connect-offsets store source offsets for source connectors

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 81

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 86

An ecommerce website maintains two topics - a high volume "purchase" topic with 5 partitions and low volume "customer" topic with 3 partitions. You would like to do a stream- table join of these topics. How should you proceed?

- A. Repartition the purchase topic to have 3 partitions
- B. Repartition customer topic to have 5 partitions
- C. Model customer as a GlobalKTable
- D. Do a KStream / KTable join after a repartition step

Answer: C

Explanation:

In case of KStream-KStream join, both need to be co-partitioned. This restriction is not applicable in case of join with GlobalKTable, which is the most efficient here.

NEW QUESTION 88

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 93

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 94

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 96

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

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 101

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 106

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

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

100% Actual & Verified — Instant Download, Please Click
[Order The CCDAK Practice Test Here](#)