Exam Questions CCDAK

Confluent Certified Developer for Apache Kafka Certification Examination

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

NEW QUESTION 1
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 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
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 ishttps://medium.com/@anyili0928/what-i-have-learned-from-kafka-partition-assignment- strategy-799fdf15d3ab

NEW QUESTION 4
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 5
You have a Kafka cluster and all the topics have a replication factor of 3. One intern at your company stopped a broker, and accidentally deleted all the data of that broker on the disk. What will happen if the broker is restarted?

A. The broker will start, and other topics will also be deleted as the broker data on the disk got deleted
B. The broker will start, and won't be online until all the data it needs to have is replicated from other leaders
C. The broker will crash
D. The broker will start, and won't have any dat
E. If the broker comes leader, we have a data loss

Answer: B

Explanation:
Kafka replication mechanism makes it resilient to the scenarios where the broker lose data on disk, but can recover from replicating from other brokers. This makes Kafka amazing!

NEW QUESTION 6
Kafka is configured with following parameters - log.retention.hours = 168 log.retention.minutes = 168 log.retention.ms = 168 How long will the messages be retained for?

A. Broker will not start due to bad configuration

B. 168 ms
C. 168 hours
D. 168 minutes

Answer: B

Explanation:
If more than one similar config is specified, the smaller unit size will take precedence.

NEW QUESTION 7
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 8
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 9
When using plain JSON data with Connect, you see the following error messageorg.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 10
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:
Seehttps://kafka.apache.org/20/documentation/streams/developer-guide/dsl-api.html#stateless-transformations

NEW QUESTION 10
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 12
An ecommerce wesbite sells some custom made goods. What's the natural way of modeling this data in Kafka streams?

A. Purchase as stream, Product as stream, Customer as stream
B. Purchase as stream, Product as table, Customer as table
C. Purchase as table, Product as table, Customer as table
D. Purchase as stream, Product as table, Customer as stream

Answer: B

Explanation:
Mostly-static data is modeled as a table whereas business transactions should be modeled as a stream.

NEW QUESTION 14
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 19
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 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 23
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 24
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 27
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 30
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 35
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 36
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 topi
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 38
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 41
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:
Seehttps://docs.confluent.io/current/streams/developer-guide/dsl-api.html#joining

NEW QUESTION 43
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 47
There are 3 producers writing to a topic with 5 partitions. There are 5 consumers consuming from the topic. How many Controllers will be present in the cluster?

A. 3
B. 5
C. 2
D. 1

Answer: D

Explanation:
There is only one controller in a cluster at all times.

NEW QUESTION 50
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 produce request with acks=all is sent to broker?

A. NotEnoughReplicasException will be returned
B. Produce request is honored with single in-sync replica
C. Produce request will block till one of the two unavailable partition is available again.

Answer: A

Explanation:
With this configuration, a single in-sync replica becomes read-only. Produce request will receive NotEnoughReplicasException.

NEW QUESTION 54
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 55
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 60
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 65
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 66
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 71
You want to sink data from a Kafka topic to S3 using Kafka Connect. There are 10 brokers in the cluster, the topic has 2 partitions with replication factor of 3. How many tasks will you configure for the S3 connector?

A. 10
B. 6
C. 3
D. 2

Answer: D

Explanation:
You cannot have more sink tasks (= consumers) than the number of partitions, so 2.

NEW QUESTION 72
.....

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