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 spreads 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
You want to perform table lookups against a KTable everytime a new record is received from the KStream. What is the output of KStream-KTable join?

A. KTable
B. GlobalKTable
C. You choose between KStream or KTable
D. Kstream

Answer: D

Explanation:
Here KStream is being processed to create another KStream.

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
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 5
A bank uses a Kafka cluster for credit card payments. What should be the value of the property unclean.leader.election.enable?

A. FALSE
B. TRUE

Answer: A

Explanation:
Setting unclean.leader.election.enable to true means we allow out-of-sync replicas to become leaders, we will lose messages when this occurs, effectively losing credit card payments and making our customers very angry.

NEW QUESTION 6
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 7
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 8
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 9
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 10
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 10
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 12
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 17
A Zookeeper ensemble contains 3 servers. Over which ports the members of the ensemble should be able to communicate in default configuration? (select three)

A. 2181
B. 3888
C. 443
D. 2888
E. 9092
F. 80

Answer: ABD

Explanation:
2181 - client port, 2888 - peer port, 3888 - leader port

NEW QUESTION 18
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
What is a generic unique id that I can use for messages I receive from a consumer?

A. topic + partition + timestamp
B. topic + partition + offset
C. topic + timestamp

Answer: B

Explanation:
(Topic,Partition,Offset) uniquely identifies a message in Kafka

NEW QUESTION 28
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 30
You are running a Kafka Streams application in a Docker container managed by Kubernetes, and upon application restart, it takes a long time for the docker container to replicate the state and get back to processing the data. How can you improve dramatically the application restart?

A. Mount a persistent volume for your RocksDB
B. Increase the number of partitions in your inputs topic
C. Reduce the Streams caching property
D. Increase the number of Streams threads

Answer: A

Explanation:
Although any Kafka Streams application is stateless as the state is stored in Kafka, it can take a while and lots of resources to recover the state from Kafka. In order to speed up recovery, it is advised to store the Kafka Streams state on a persistent volume, so that only the missing part of the state needs to be recovered.

NEW QUESTION 31
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 35
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 38
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 39
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 42
A consumer wants to read messages from a specific partition of a topic. How can this be achieved?

A. Call subscribe(String topic, int partition) passing the topic and partition number as the arguments
B. Call assign() passing a Collection of TopicPartitions as the argument
C. Call subscribe() passing TopicPartition as the argument

Answer: B

Explanation:
assign() can be used for manual assignment of a partition to a consumer, in which case subscribe() must not be used. Assign() takes a collection of TopicPartition object as an argument https://kafka.apache.org/23/javadoc/org/apache/kafka/clients/consumer/KafkaConsumer.ht
ml#assign-java.util.Collection-

NEW QUESTION 44
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 47
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 49
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 52
Which of the following event processing application is stateless? (select two)

A. Read events from a stream and modifies them from JSON to Avro
B. Publish the top 10 stocks each day
C. Read log messages from a stream and writes ERROR events into a high-priority stream and the rest of the events into a low-priority stream
D. Find the minimum and maximum stock prices for each day of trading

Answer: AC

Explanation:
Stateless means processing of each message depends only on the message, so converting from JSON to Avro or filtering a stream are both stateless operations

NEW QUESTION 55
Partition leader election is done by

A. The consumers
B. The Kafka Broker that is the Controller
C. Zookeeper
D. Vote amongst the brokers

Answer: C

Explanation:
The Controller is a broker that is responsible for electing partition leaders

NEW QUESTION 57
How do Kafka brokers ensure great performance between the producers and consumers? (select two)

A. It compresses the messages as it writes to the disk
B. It leverages zero-copy optimisations to send data straight from the page-cache
C. It buffers the messages on disk, and sends messages from the disk reads
D. It transforms the messages into a binary format
E. It does not transform the messages

Answer: BE

Explanation:
Kafka transfers data with zero-copy and sends the raw bytes it receives from the producer
straight to the consumer, leveraging the RAM available as page cache

NEW QUESTION 60
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 61
How will you find out all the partitions without a leader?

A. kafka-topics.sh --broker-list localhost:9092 --describe --under-replicated-partitions
B. kafka-topics.sh --bootstrap-server localhost:2181 --describe --unavailable-partitions
C. kafka-topics.sh --zookeeper localhost:2181 --describe --unavailable-partitions
D. kafka-topics.sh --zookeeper localhost:2181 --describe --under-replicated-partitions

Answer: C

Explanation:
Please note that as of Kafka 2.2, the --zookeeper option is deprecated and you can now usekafka-topics.sh --bootstrap-server localhost:9092 --describe --unavailable-partitions

NEW QUESTION 62
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 settingenable.idempotence=true which takes care of the out of ordering case on its own.
Seehttps://issues.apache.org/jira/browse/KAFKA-5494

NEW QUESTION 66
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 68
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 72
You are sending messages with keys to a topic. To increase throughput, you decide to increase the number of partitions of the topic. Select all that apply.

A. All the existing records will get rebalanced among the partitions to balance load
B. New records with the same key will get written to the partition where old records with that key were written
C. New records may get written to a different partition
D. Old records will stay in their partitions

Answer: CD

Explanation:
Increasing the number of partition causes new messages keys to get hashed differently, and breaks the guarantee "same keys goes to the same partition". Kafka logs are immutable and the previous messages are not re-shuffled

NEW QUESTION 73
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 76
The exactly once guarantee in the Kafka Streams is for which flow of data?

A. Kafka => Kafka
B. Kafka => External
C. External => Kafka

Answer: A

Explanation:
Kafka Streams can only guarantee exactly once processing if you have a Kafka to Kafka topology.

NEW QUESTION 81
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 86
.....

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