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
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
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 4
In Kafka Streams, by what value are internal topics prefixed by?

A. tasks-<number>
B. application.id
C. group.id
D. kafka-streams-

Answer: B

Explanation:
In Kafka Streams, the application.id is also the underlying group.id for your consumers, and the prefix for all internal topics (repartition and state)

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
Where are the dynamic configurations for a topic stored?

A. In Zookeeper

B. In an internal Kafka topic topic_configuratins
C. In server.properties
D. On the Kafka broker file system

Answer: A

Explanation:
Dynamic topic configurations are maintained in Zookeeper.

NEW QUESTION 7
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 8
How will you set the retention for the topic named ‚Ä??my-topic‚Ä?? to 1 hour?

A. Set the broker config log.retention.ms to 3600000
B. Set the consumer config retention.ms to 3600000
C. Set the topic config retention.ms to 3600000
D. Set the producer config retention.ms to 3600000

Answer: C

Explanation:
retention.ms can be configured at topic level while creating topic or by altering topic. It shouldn't be set at the broker level (log.retention.ms) as this would impact all the topics in the cluster, not just the one we are interested in

NEW QUESTION 9
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 10
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 10
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 13
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 14
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 18
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 20
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 25
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 28
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 32
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 35
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 37
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 38
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 42
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 44

```
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 47
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 52
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 57
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 60
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 64
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 68
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 72
To get acknowledgement of writes to only the leader partition, we need to use the config...

A. acks=1
B. acks=0
C. acks=all

Answer: A

Explanation:
Producers can set acks=1 to get acknowledgement from partition leader only.

NEW QUESTION 74
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 77
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 80
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 82
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.html#assign-java.util.Collection-

NEW QUESTION 87
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 91
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 92
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 93
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 96
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 99
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 103
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 108
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 110
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 114
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 117
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 122
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 123
What happens when broker.rack configuration is provided in broker configuration in Kafka cluster?

A. You can use the same broker.id as long as they have different broker.rack configuration
B. Replicas for a partition are placed in the same rack
C. Replicas for a partition are spread across different racks
D. Each rack contains all the topics and partitions, effectively making Kafka highly available

Answer: C

Explanation:
Partitions for newly created topics are assigned in a rack alternating manner, this is the only change broker.rack does

NEW QUESTION 125
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 127
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 131
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 132
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 137
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 139
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 141
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 144
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 148
What information isn't stored inside of Zookeeper? (select two)

A. Schema Registry schemas
B. Consumer offset
C. ACL inforomation
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 150
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 154
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 159
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 164
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 165
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 paritions increased.

NEW QUESTION 166
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 170
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 172
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 173
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 174
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 176
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 178
To allow consumers in a group to resume at the previously committed offset, I need to set the proper value for...

A. value.deserializer
B. auto.offset.resets
C. group.id
D. enable.auto.commit

Answer: C

Explanation:
Setting a group.id that's consistent across restarts will allow your consumers part of the same group to resume reading from where offsets were last committed for that group

NEW QUESTION 183
What is not a valid authentication mechanism in Kafka?

A. SASL/GSSAPI
B. SASL/SCRAM
C. SAML
D. SSL

Answer: C

Explanation:
Learn more about security herehttps://kafka.apache.org/documentation/#security

NEW QUESTION 186
The rule "same key goes to the same partition" is true unless...

A. the number of producer changes
B. the number of kafka broker changes
C. the number of partition changes
D. the replication factor changes

Answer: C

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 187
Which of the following errors are retriable 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 retriable errors, others non-retriable errors. See the full list of errors and their "retriable" status herehttps://kafka.apache.org/protocol#protocol_error_codes

NEW QUESTION 190
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 192
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 194
.....

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