Confluent

Exam Questions CCDAK

Confluent Certified Developer for Apache Kafka Certification Examination

NEW QUESTION 1
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 2
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 3
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 4
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 5
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 6
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 7
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 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 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
There are 3 brokers in the cluster. You want to create a topic with a single partition that is resilient to one broker failure and one broker maintenance. What is the replication factor will you specify while creating the topic?

A. 6
B. 3
C. 2
D. 1

Answer: B

Explanation:
1 is not possible as it doesn't provide resilience to failure, 2 is not enough as if we take a broker down for maintenance, we cannot tolerate a broker failure, and 6 is impossible as we only have 3 brokers (RF cannot be greater than the number of brokers). Here the correct answer is 3

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 11
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 13
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 18
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 23
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 26
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 30
Your streams application is reading from an input topic that has 5 partitions. You run 5 instances of your application, each with num.streams.threads set to 5. How many stream tasks will be created and how many will be active?

A. 5 created, 1 active
B. 5 created, 5 active
C. 25 created, 25 active
D. 25 created, 5 active

Answer: D

Explanation:
One partition is assigned a thread, so only 5 will be active, and 25 threads (i.e. tasks) will be created

NEW QUESTION 33
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 35
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 40
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 43
A Kafka producer application wants to send log messages to a topic that does not include any key. What are the properties that are mandatory to configure for the producer configuration? (select three)

A. bootstrap.servers
B. partition
C. key.serializer
D. value.serializer
E. key
F. value

Answer: ACD

Explanation:
Both key and value serializer are mandatory.

NEW QUESTION 45
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 50
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 55
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 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 61
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 63
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 67
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 69
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 71
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 73
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 75
A producer application in a developer machine was able to send messages to a Kafka topic. After copying the producer application into another developer's machine, the producer is able to connect to Kafka but unable to produce to the same Kafka topic because of an authorization issue. What is the likely issue?

A. Broker configuration needs to be changed to allow a different producer
B. You cannot copy a producer application from one machine to another
C. The Kafka ACL does not allow another machine IP
D. The Kafka Broker needs to be rebooted

Answer: C

Explanation:
ACLs take "Host" as a parameter, which represents an IP. It can be * (all IP), or a specific IP. Here, it's a specific IP as moving a producer to a different machine breaks the consumer, so the ACL needs to be updated

NEW QUESTION 79
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 84
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 88
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 92
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 96
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 101
Using the Confluent Schema Registry, where are Avro schema stored?

A. In the Schema Registry embedded SQL database
B. In the Zookeeper node /schemas
C. In the message bytes themselves
D. In the _schemas topic

Answer: D

Explanation:
The Schema Registry stores all the schemas in the _schemas Kafka topic

NEW QUESTION 106
A topic has three replicas and you set min.insync.replicas to 2. If two out of three replicas are not available, what happens when a consume request is sent to broker?

A. Data will be returned from the remaining in-sync replica

B. An empty message will be returned
C. NotEnoughReplicasException will be returned
D. A new leader for the partition will be elected

Answer: A

Explanation:
With this configuration, a single in-sync replica is still readable, but not writeable if the producer using acks=all

NEW QUESTION 108
Which of the following is true regarding thread safety in the Java Kafka Clients?

A. One Producer can be safely used in multiple threads
B. One Consumer can be safely used in multiple threads
C. One Consumer needs to run in one thread
D. One Producer needs to be run in one thread

Answer: AC

Explanation:
KafkaConsumer is not thread-safe, KafkaProducer is thread safe.

NEW QUESTION 111
How often is log compaction evaluated?

A. Every time a new partition is created
B. Every time a segment is closed
C. Every time a message is sent to Kafka
D. Every time a message is flushed to disk

Answer: B

Explanation:
Log compaction is evaluated every time a segment is closed. It will be triggered if enough data is "dirty" (see dirty ratio config)

NEW QUESTION 113
.....

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