



Vendor: Confluent

Exam Code: CCDAK

Exam Name: Confluent Certified Developer for Apache
Kafka

Version: DEMO

QUESTION 1

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.

QUESTION 2

Which is an optional field in an Avro record?

- A. doc
- B. name
- C. namespace
- D. fields

Answer: A

QUESTION 3

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)

QUESTION 4

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.

QUESTION 5

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.

QUESTION 6

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

QUESTION 7

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.

QUESTION 8

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

QUESTION 9

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

QUESTION 10

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.

QUESTION 11

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.

QUESTION 12

To enhance compression, I can increase the chances of batching by using

- A. acks=all
- B. linger.ms=20
- C. batch.size=65536
- D. max.message.size=10MB

Answer: B

Explanation:

linger.ms forces the producer to wait before sending messages, hence increasing the chance of creating batches that can be heavily compressed.

QUESTION 13

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.

QUESTION 14

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.

QUESTION 15

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()

QUESTION 16

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).

QUESTION 17

Which of the following Kafka Streams operators are stateful? (select all that apply)

- A. flatmap
- B. reduce
- C. joining
- D. count
- E. peek
- F. aggregate

Answer: BCDF

QUESTION 18

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.

QUESTION 19

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. 0

- C. 2
- D. 1

Answer: C

Explanation:

Two brokers can go down, and one replica will still be able to receive and serve data.

QUESTION 20

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.

QUESTION 21

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. This cannot be changed.
- C. 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'`;

QUESTION 22

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.

QUESTION 23

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.

QUESTION 24

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.

QUESTION 25

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.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.lead2pass.com/all-products.html>



Microsoft



ORACLE



JUNIPER
NETWORKS



EMC²
where information lives

10% Discount Coupon Code: ASTR14