Vendor: Confluent

Exam Code: CCDAK

Exam Name: Confluent Certified Developer for Apache Kafka

Version: DEMO

QUESTION 1
You need to consume messages from Kafka using the command-line interface (CLI). Which command should you use?

A. kafka-console-consumer
B. kafka-consumer
C. kafka-get-message
D. kafka- consume

Answer: A
Explanation:
kafka-console-consumer is the correct CLI tool for consuming messages from Kafka topics. It allows you to read messages from a Kafka cluster directly from the command line.

QUESTION 2
You are experiencing low throughput from a Java producer. You monitor the Kafka metrics and notice that the metrics for that producer show low I/O thread ratio and low I/O thread wait ratio. What is the most likely the cause of the slow producer performance?

A. Compression is enabled.
B. The producer is sending large batches of messages.
C. There is a bad data link layer (layer 2) connection from the producer to the cluster.
D. The producer code has an expensive callback function.

Answer: D
Explanation:
A low I/O thread ratio and low I/O thread wait ratio indicate that the producer's I/O threads are busy and not waiting on network or disk I/O. This suggests the bottleneck is in the producer's application code, such as an expensive callback function, which slows down processing.

QUESTION 3
Which configuration is valid for deploying a JDBC Source Connector to read all the rows from orders table and write it to db1-orders topic?

A.
```
{
"name":"jdbc-source",
"connector.class": "io.confluent.connect.jdbc.JdbcSourceConnector",
"tasks.max": "1",
"connection.url": "jdbc:mysql://mysql:3306/db1?user=user&password=password&useSSL=false",
"topic.prefix": "db1-",
"table.whitelist": "orders"
}
```
B.
```
{
"name":"db1-orders".
"connector.class": "io.confluent.connect.jdbc.JdbcSourceConnector",
"tasks.max": "1",
"connection.url": "jdbc:mysql://mysql:3306/db1?user=user&password=password&useSSL=false",
"topic.prefix": "db1-",
"table.blacklist": "ord*"
}
```
C.
```
{
"name":"orders-connect",
```

```
"connector.class": "io.confluent.connect.jdbc.JdbcSourceConnector",
"tasks.max": "1",
"connection.url": "jdbc:mysql://mysql:3306/db1",
"topic.whitelist": "orders",
"auto.create": "true"
}
```

D.
```
{
"name":"jdbc-source",
"connector.class": "io.confluent.connect.jdbc.JdbcSourceConnector",
"tasks.max": "1",
"connection.url": "jdbc:mysql://mysql:3306/db1?user=user&useAutoAuthentication=true",
"topic.prefix": "db1-",
"table.whitelist": "orders"
}
```

Answer: A
Explanation:
Chosen option correctly configures the JdbcSourceConnector with a valid connection.url, specifies table.whitelist to read from the orders table, and sets topic.prefix so that data from the orders table will be written to the db1-orders topic. The syntax and required fields are correct for this use case.

QUESTION 4
You are developing a Kafka Streams application with a complex topology that has multiple sources, processors, sinks, and sub-topologies. You are working in a development environment and do not have access to a real Kafka Cluster or Topic.

You need to perform Unit Test on your Kafka Streams application.

Which should you use?

A. TopologyTestDriver
B. MockProducer, MockConsumer
C. TestProducer, TestConsumer
D. KafkaUnitTestDriver

Answer: A
Explanation:
TopologyTestDriver is specifically designed for unit testing Kafka Streams applications. It allows you to simulate input and output to your stream topology without needing a running Kafka cluster or real topics.

QUESTION 5
Which two statements are correct about transactions in Kafka? (Choose two.)

A. All messages from a failed transaction will be deleted from a Kafka topic.
B. Transactions are only possible when writing messages to a topic with single partition.
C. Consumers can consume both committed and uncommitted transactions.
D. Information about producers and their transactions is stored in the '__transaction_state' topic.
E. Transactions guarantee at least once delivery of messages.

Answer: CD

Explanation:
Consumers can consume both committed and uncommitted transactions, but by default, only committed messages are visible to consumers with isolation.level=read_committed.

Information about producers and their transactions is stored in the __transaction_state topic, which is managed internally by Kafka for transaction coordination.

QUESTION 6
Drag and Drop Question

Match the topic configuration setting with the reason the setting affects topic durability.

To answer, move each of the options below to the corresponding answer area. Partial credit is given for each correct answer.

Options	Answer Area	
Prevents data loss by only considering in-sync replicas when rebalancing.	unclean.leader.election.enable=false	
Specifies how many redundant copies of partitions that are distributed across the brokers.	Min.insync.replicas=2	
Sets the standard for the number of partition instances that must keep up with the latest committed message.	replication.factor=3	

Answer:

Options	Answer Area	
Prevents data loss by only considering in-sync replicas when rebalancing.	unclean.leader.election.enable=false	Prevents data loss by only considering in-sync replicas when rebalancing.
Specifies how many redundant copies of partitions that are distributed across the brokers.	Min.insync.replicas=2	Sets the standard for the number of partition instances that must keep up with the latest committed message.
Sets the standard for the number of partition instances that must keep up with the latest committed message.	replication.factor=3	Specifies how many redundant copies of partitions that are distributed across the brokers.

QUESTION 7
You are designing a stream pipeline to monitor the real-time location of GPS trackers, where historical location data is not required.

Each event in the stream has the following structure:

```
- Key: tracker_id
- Value: latitude, longitude
```

You need to ensure that the latest location for each tracker is always retained in the Kafka topic.

Which topic configuration parameter should you set?

A. retention.ms = 0
B. min.cleanable.dirty.ratio = -1
C. cleanup.policy = compact
D. retention.ms = infinite

Answer: C
Explanation:
Setting cleanup.policy = compact enables log compaction, ensuring that only the latest value for each key (tracker_id) is retained in the topic, which is ideal for storing the most recent GPS location per tracker.

QUESTION 8
You need to collect logs from a host and write them to a Kafka topic named 'logs-topic'. You decide to use Kafka Connect File Source connector for this task. What is the preferred deployment mode for this connector?

A. Standalone mode
B. Distributed mode
C. Parallel mode
D. SingleCluster mode

Answer: A
Explanation:
Standalone mode is preferred for simple use cases like collecting logs from a single host with the Kafka Connect File Source connector. It is easy to set up and manage for local or single-node data ingestion.

QUESTION 9
You create a topic named 'IoT-Data' with 10 partitions and replication factor of three. A producer application is producing messages with 1 mb size. The messages need to be compressed using Gzip compression. Which two statements are true in this scenario? (Choose two.)

A. Compression type will be stored in batch attributes.
B. By default, compression is the producer's responsibility.
C. The message is already compressed so it will not be serialized.
D. All compressed messages will be stored in the same topic partition.

Answer: AB
Explanation:
The compression type is stored in the batch attributes, allowing consumers to know how to decompress the data.

By default, message compression is handled by the producer, and the broker stores the messages in their compressed form.

QUESTION 10
The producer code below features a 'Callback' class with a method called 'onCompletion()'.

```
class MyCallback implements Callback {
   ProducerRecord<String, String> messageString;

   public MyCallback(ProducerRecord<String, String> messageString) {
     this.messageString = messageString;
   }

   public void onCompletion(RecordMetadata metadata, Exception exception) {
     if (exception != null) {
       e.printStackTrace();
     } else {
       System.out.println("Message String = " + messageString.value() + ", Offset = " + metadata.offset());
     }
   }
}

producer.send(record, new MyCallback(record));
```

When will the 'onCompletion()' method be invoked?

A. When a consumer sends an acknowledgement to the producer
B. When the producer puts the message into its socket buffer
C. When the producer batches the message
D. When the producer receives the acknowledgment from the broker

Answer: D
Explanation:
The onCompletion() method of the Callback is invoked when the producer receives the acknowledgment from the broker, indicating that the message has been successfully written or an error has occurred.

QUESTION 11
What is the purpose of the key.serializer property?

A. key.serializer specifies the serializer class to use for the key of the message.
B. key.serializer specifies the serializer class to use for the value of the message.
C. key.serializer determines which class to use for the message.
D. key.serializer determines where to send the message.

Answer: A
Explanation:
The key.serializer property specifies the serializer class to use for the key of the message when producing data to Kafka.

QUESTION 12
You are working on an Orders microservice that is publishing messages to an Orders topic. It is a business requirement that the Orders for a given customer are processed sequentially. The Orders topic is partitioned for scalability.

You need to ensure the business requirements are met during application development.

Which action should you take?

A. Enable producer Idempotence. Ensure the message key is set to current timestamp, set batch.size =1, and set linger.ms to 0, and the number of partitions is not changed.
B. Enable producer Idempotence. Ensure the message key is set to customer id, set max-in-flight requests per connection to less than or equal to 5, and the number of partitions is not changed.
C. Enable producer Idempotence. Ensure the message key is set to producer id, set max-in-flight requests per connection to 1, and the number of partitions is changed.
D. Enable producer Idempotence. Ensure the message key is set to current timestamp, set linger.ms to 0, set acks=0, and the number of partitions is not changed.

Answer: B
Explanation:
Setting the message key to the customer id ensures that all orders for a given customer go to the same partition and are processed sequentially. Enabling idempotence and keeping max-in-flight requests per connection to ≤5 ensures ordering and exactly-once delivery, while not changing the number of partitions maintains key-to-partition mapping.

QUESTION 13
Which configuration determines how many bytes of data are collected before sending messages to the Kafka broker?

A. batch.size
B. max.block.size
C. buffer.memory
D. send.buffer.bytes

Answer: A
Explanation:
The batch.size configuration determines how many bytes of data are collected before the producer sends messages in a batch to the Kafka broker.

QUESTION 14
Your company has three Kafka clusters: Development, Testing, and Production. The Production cluster is running out of storage so you add a new node to the Kafka cluster. Which two statements about the new node are true? (Choose two.)

A. A node id will be assigned to the new node automatically.
B. A newly added node will have KRaft controller role by default.
C. A new node will not have any partitions assigned to it unless a new topic is created.
D. A new node can be added without stopping existing cluster nodes.

Answer: CD
Explanation:
A new node will not have any partitions assigned to it unless a new topic is created or partitions are rebalanced across the cluster.

A new node can be added to a running Kafka cluster without stopping any existing nodes, supporting Kafka's horizontal scalability.

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

10% Discount Coupon Code: **ASTR14**