



Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.



CCDAK Dumps
CCDAK Braindumps
CCDAK Real Questions
CCDAK Practice Test
CCDAK Actual Questions



killexams.com

Confluent

CCDAK

Confluent Certified Developer for Apache Kafka

ORDER FULL VERSION

<https://killexams.com/pass4sure/exam-detail/CCDAK>



Question: 354

Which of the following is NOT a valid Kafka Connect connector type?

- A. Source Connector
- B. Sink Connector
- C. Processor Connector
- D. Transform Connector

Answer: C

Explanation: "Processor Connector" is not a valid Kafka Connect connector type. The valid connector types are Source Connector (for importing data into Kafka), Sink Connector (for exporting data from Kafka), and Transform Connector (for modifying or transforming data during the import or export process).

Question: 355

Which of the following is a benefit of using Apache Kafka for real-time data streaming?

- A. High-latency message delivery
- B. Centralized message storage and processing
- C. Limited scalability and throughput
- D. Inability to handle large volumes of data
- E. Fault-tolerance and high availability

Answer: E

Explanation: One of the benefits of using Apache Kafka for real-time data streaming is its fault-tolerance and high availability. Kafka is designed to provide durability, fault tolerance, and high availability of data streams. It can handle large volumes of data and offers high scalability and throughput. Kafka also allows for centralized message storage and processing, enabling real-time processing of data from multiple sources.

Question: 356

Which of the following is NOT a valid deployment option for Kafka?

- A. On-premises deployment
- B. Cloud deployment (e.g., AWS, Azure)
- C. Containerized deployment (e.g., Docker)
- D. Mobile deployment (e.g., Android, iOS)

Answer: D

Explanation: Mobile deployment (e.g., Android, iOS) is not a valid deployment option for Kafka. Kafka is typically deployed in server or cloud environments to handle high-throughput and real-time data streaming. It is commonly deployed on servers in on-premises data centers or in the cloud, such as AWS (Amazon Web Services) or Azure. Kafka can also be containerized using technologies like Docker and deployed in container orchestration platforms like Kubernetes. However, deploying Kafka on mobile platforms like Android or iOS is not a typical use case. Kafka is designed for server-side data processing and messaging, and it is not optimized for mobile devices.

Question: 357

Which of the following is a feature of Kafka Streams?

- A. It provides a distributed messaging system for real-time data processing.
- B. It supports exactly-once processing semantics for stream processing.
- C. It enables automatic scaling of Kafka clusters based on load.

Answer: B

Explanation: Kafka Streams supports exactly-once processing semantics for stream processing. This means that when processing data streams using Kafka Streams, each record is processed exactly once, ensuring data integrity and consistency. This is achieved through a combination of Kafka's transactional messaging and state management features in Kafka Streams.

Question: 358

When designing a Kafka consumer application, what is the purpose of setting the `auto.offset.reset` property?

- A. To control the maximum number of messages to be fetched per poll.
- B. To specify the topic to consume messages from.
- C. To determine the behavior when there is no initial offset in Kafka or if the current offset does not exist.
- D. To configure the maximum amount of time the consumer will wait for new messages.

Answer: C

Explanation: The `auto.offset.reset` property is used to determine the behavior when there is no initial offset in Kafka or if the current offset does not exist. It specifies whether the consumer should automatically reset the offset to the earliest or latest available offset in such cases.

Question: 359

What is the role of a Kafka producer?

- A. To consume messages from Kafka topics and process them.
- B. To store and manage the data in Kafka topics.
- C. To replicate Kafka topic data across multiple brokers.
- D. To publish messages to Kafka topics.

Answer: D

Explanation: The role of a Kafka producer is to publish messages to Kafka topics. Producers are responsible for sending messages to Kafka brokers, which then distribute the messages to the appropriate partitions of the specified topics. Producers can be used to publish data in real-time or batch mode to Kafka for further processing or consumption.

Question: 360

Which of the following is a valid way to configure Kafka producer retries?

- A. Using the `retries` property in the producer configuration

- B. Using the `retry.count` property in the producer configuration
- C. Using the `producer.retries` property in the producer configuration
- D. Using the `producer.retry.count` property in the producer configuration

Answer: A

Explanation: Kafka producer retries can be configured using the `retries` property in the producer configuration. This property specifies the number of retries that the producer will attempt in case of transient failures.

Question: 361

Which of the following is NOT a valid approach for Kafka cluster scalability?

- A. Increasing the number of brokers
- B. Increasing the number of partitions per topic
- C. Increasing the replication factor for topics
- D. Increasing the retention period for messages

Answer: D

Explanation: Increasing the retention period for messages is not a valid approach for Kafka cluster scalability. The retention period determines how long messages are retained within Kafka, but it does not directly impact the scalability of the cluster. Valid approaches for scalability include increasing the number of brokers, partitions, and replication factor.

Question: 362

Which of the following is NOT a core component of Apache Kafka?

- A. ZooKeeper
- B. Kafka Connect
- C. Kafka Streams
- D. Kafka Manager

Answer: D

Explanation: ZooKeeper, Kafka Connect, and Kafka Streams are all core components of Apache Kafka. ZooKeeper is used for coordination, synchronization, and configuration management in Kafka. Kafka Connect is a framework for connecting Kafka with external systems. Kafka Streams is a library for building stream processing applications with Kafka. However, "Kafka Manager" is not a core component of Kafka. It is a third-party tool used for managing and monitoring Kafka clusters.

Question: 363

Which of the following is true about Kafka replication?

- A. Kafka replication ensures that each message in a topic is stored on multiple brokers for fault tolerance.
- B. Kafka replication is only applicable to log-compacted topics.
- C. Kafka replication allows data to be synchronized between Kafka and external systems.
- D. Kafka replication enables compression and encryption of messages in Kafka.

Answer: A

Explanation: Kafka replication ensures fault tolerance by storing multiple

copies of each message in a topic across different Kafka brokers. Each topic partition can have multiple replicas, and Kafka automatically handles replication and leader election to ensure high availability and durability of data.

Question: 364

What is Kafka log compaction?

- A. A process that compresses the Kafka log files to save disk space.
- B. A process that removes duplicate messages from Kafka topics.
- C. A process that deletes old messages from Kafka topics to free up disk space.
- D. A process that retains only the latest value for each key in a Kafka topic.

Answer: D

Explanation: Kafka log compaction is a process that retains only the latest value for each key in a Kafka topic. It ensures that the log maintains a compact representation of the data, removing any duplicate or obsolete messages. Log compaction is useful when the retention of the full message history is not required, and only the latest state for each key is needed.

Question: 365

What is the significance of the acks configuration parameter in the Kafka producer?

- A. It determines the number of acknowledgments the leader broker must receive before considering a message as committed.
- B. It defines the number of replicas that must acknowledge the message before

considering it as committed.

C. It specifies the number of retries the producer will attempt in case of failures before giving up.

D. It sets the maximum size of messages that the producer can send to the broker.

Answer: A

Explanation: The acks configuration parameter in the Kafka producer determines the number of acknowledgments the leader broker must receive before considering a message as committed. It can be set to "all" (which means all in-sync replicas must acknowledge), "1" (which means only the leader must acknowledge), or a specific number of acknowledgments.

Question: 366

Which of the following is NOT a valid method for handling Kafka message serialization?

A. JSON

B. Avro

C. Protobuf

D. XML

Answer: D

Explanation: "XML" is not a valid method for handling Kafka message serialization. Kafka supports various serialization formats such as JSON, Avro, and Protobuf, but not XML.

Question: 367

Which of the following is the correct command to create a new consumer group in Apache Kafka?

- A. `kafka-consumer-groups.sh --bootstrap-server localhost:9092 --create --group my_group`
- B. `kafka-consumer-groups.sh --create --group my_group`
- C. `kafka-consumer-groups.sh --bootstrap-server localhost:2181 --create --group my_group`
- D. `kafka-consumer-groups.sh --group my_group --create`

Answer: A

Explanation: The correct command to create a new consumer group in Apache Kafka is "`kafka-consumer-groups.sh --bootstrap-server localhost:9092 --create --group my_group`". This command creates a new consumer group with the specified group name. The "`--bootstrap-server`" option specifies the Kafka bootstrap server, and the "`--group`" option specifies the consumer group name. The other options mentioned either have incorrect parameters or do not include the necessary bootstrap server information.

Question: 368

What is the purpose of a Kafka producer in Apache Kafka?

- A. To consume messages from Kafka topics
- B. To manage the replication of data across Kafka brokers
- C. To provide fault tolerance by distributing the load across multiple consumers
- D. To publish messages to Kafka topics

Answer: D

Explanation: The purpose of a Kafka producer in Apache Kafka is to publish messages to Kafka topics. Producers are responsible for creating and sending messages to Kafka brokers, which then distribute the messages to the appropriate partitions of the topics. Producers can specify the topic and partition to which a message should be sent, as well as the key and value of the message. They play a crucial role in the data flow of Kafka by publishing new messages for consumption by consumers.

Question: 369

What is the purpose of the Kafka Connect Transformer?

- A. To convert Kafka messages from one topic to another
- B. To transform the data format of Kafka messages
- C. To perform real-time stream processing within a Kafka cluster
- D. To manage and monitor the health of Kafka Connect connectors

Answer: B

Explanation: The Kafka Connect Transformer is used to transform the data format of Kafka messages during the import or export process. It allows for the modification, enrichment, or restructuring of the data being transferred between Kafka and external systems by applying custom transformations to the messages.

SAMPLE QUESTIONS



*These questions are for demo purpose only. **Full version** is up to date and contains actual questions and answers.*

Killexams.com is an online platform that offers a wide range of services related to certification exam preparation. The platform provides actual questions, exam dumps, and practice tests to help individuals prepare for various certification exams with confidence. Here are some key features and services offered by Killexams.com:



Actual Exam Questions: *Killexams.com provides actual exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these actual questions, candidates can familiarize themselves with the content and format of the real exam.*

Exam Dumps: *Killexams.com offers exam dumps in PDF format. These dumps contain a comprehensive collection of questions and answers that cover the exam topics. By using these dumps, candidates can enhance their knowledge and improve their chances of success in the certification exam.*

Practice Tests: *Killexams.com provides practice tests through their desktop VCE exam simulator and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice tests cover a wide range of questions and enable candidates to identify their strengths and weaknesses.*

Guaranteed Success: *Killexams.com offers a success guarantee with their exam dumps. They claim that by using their materials, candidates will pass their exams on the first attempt or they will refund the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exams.*

Updated Content: *Killexams.com regularly updates its question bank and exam dumps to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.*

Technical Support: *Killexams.com provides free 24x7 technical support to assist candidates with any queries or issues they may encounter while using their services. Their certified experts are available to provide guidance and help candidates throughout their exam preparation journey.*

For More exams visit <https://killexams.com/vendors-exam-list>
Kill your exam at First Attempt....Guaranteed!