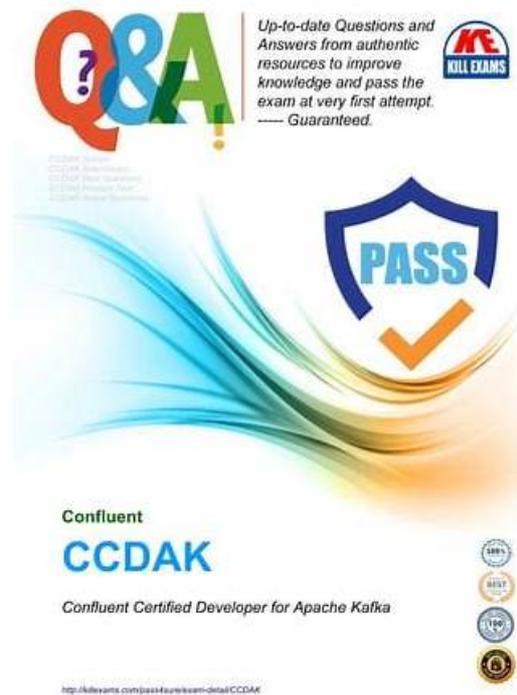


# Latest Braindumps Confluent CCAAK Ebook - CCAAK Latest Exam Pattern



BTW, DOWNLOAD part of DumpsReview CCAAK dumps from Cloud Storage: [https://drive.google.com/open?id=1EgfXUcTPmT6M8i6BpPHtlnrHW\\_AXfXbM](https://drive.google.com/open?id=1EgfXUcTPmT6M8i6BpPHtlnrHW_AXfXbM)

Our App online version of CCAAK study materials, it is developed on the basis of a web browser, as long as the user terminals on the browser, can realize the application which has applied by the CCAAK simulating materials of this learning model, users only need to open the App link, you can quickly open the learning content in real time in the ways of the CCAAK Exam Guide, can let users anytime, anywhere learning through our App, greatly improving the use value of our CCAAK exam prep.

## Confluent CCAAK Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"> <li>• <b>Troubleshooting:</b> This section of the exam measures skills of a Kafka Administrator and includes diagnosing common issues in Kafka clusters. It covers problem areas such as performance bottlenecks, message delivery failures, replication issues, and consumer lag, along with techniques to resolve them effectively.</li> </ul>
Topic 2	<ul style="list-style-type: none"> <li>• <b>Apache Kafka® Fundamentals:</b> This section of the exam measures skills of a Kafka Administrator and covers core concepts such as Kafka architecture, components, and data flow. It assesses the candidate's understanding of topics like topics and partitions, brokers, producers, consumers, and message retention.</li> </ul>

Topic 3	<ul style="list-style-type: none"> <li>• <b>Deployment Architecture:</b> This section of the exam measures skills of a Kafka Administrator and covers different Kafka deployment topologies. It evaluates understanding of high availability, disaster recovery, multi-data center replication, and the roles of Zookeeper or KRaft in the cluster architecture.</li> </ul>
Topic 4	<ul style="list-style-type: none"> <li>• <b>Observability:</b> This section of the exam measures skills of a Site Reliability Engineer and focuses on monitoring Kafka clusters. It assesses knowledge of metrics, logging, and alerting tools, including how to use them to maintain cluster health and performance visibility.</li> </ul>
Topic 5	<ul style="list-style-type: none"> <li>• <b>Apache Kafka® Cluster Configuration:</b> This section of the exam measures skills of a Kafka Administrator and includes configuring broker properties, tuning for performance, managing topic-level settings, and applying best practices for production-grade environments.</li> </ul>
Topic 6	<ul style="list-style-type: none"> <li>• <b>Kafka Connect:</b> This section of the exam measures skills of a Site Reliability Engineer and addresses the use and management of Kafka Connect for data integration. It includes setting up connectors, managing configurations, and ensuring efficient movement of data between Kafka and external systems.</li> </ul>

>> Latest Braindumps Confluent CCAAK Ebook <<

## 100% Pass Quiz 2026 Pass-Sure Confluent CCAAK: Latest Braindumps Confluent Certified Administrator for Apache Kafka Ebook

But the helpful feature is that it works without a stable internet service. What makes your Confluent Certification Exams preparation super easy is it imitates the exact syllabus and structure of the actual Confluent CCAAK Certification Exam. DumpsReview never leaves its customers in the lurch.

### Confluent Certified Administrator for Apache Kafka Sample Questions (Q31-Q36):

#### NEW QUESTION # 31

A company is setting up a log ingestion use case where they will consume logs from numerous systems. The company wants to tune Kafka for the utmost throughput.

In this scenario, what acknowledgment setting makes the most sense?

- A. acks=1
- B. acks=0
- C. acks=undefined
- D. acks=all

**Answer: B**

Explanation:

acks=0 provides the highest throughput because the producer does not wait for any acknowledgment from the broker. This minimizes latency and maximizes performance.

However, it comes at the cost of no durability guarantees - messages may be lost if the broker fails before writing them. This setting is suitable when throughput is critical and occasional data loss is acceptable, such as in some log ingestion use cases where logs are also stored elsewhere.

#### NEW QUESTION # 32

Per customer business requirements, a system's high availability is more important than message reliability.

Which of the following should be set?

- A. The number of brokers in the cluster should be always odd (3, 5, 7 and so on).
- B. The linger.ms should be set to '0'.
- C. Message retention.ms should be set to -1.
- D. Unclean leader election should be enabled.

**Answer: D**

Explanation:

Enabling unclean leader election allows Kafka to elect a non-in-sync replica as leader if all in-sync replicas are unavailable. This sacrifices message reliability (possible data loss) in favor of high availability, aligning with the requirement.

### NEW QUESTION # 33

Multiple clients are sharing a Kafka cluster.

As an administrator, how would you ensure that Kafka resources are distributed fairly to all clients?

- **A. Quotas**
- B. Consumer Groups
- C. ACLs
- D. Rebalancing

**Answer: A**

Explanation:

Kafka quotas allow administrators to control and limit the rate of data production and consumption per client (producer/consumer), ensuring fair use of broker resources among multiple clients.

### NEW QUESTION # 34

You are managing a Kafka cluster with five brokers (broker id '0', '1','2','3','4') and three ZooKeepers. There are 100 topics, five partitions for each topic, and replication factor three on the cluster. Broker id '0' is currently the Controller, and this broker suddenly fails.

Which statements are correct? (Choose three.)

- A. The broker id is used as the epoch number to prevent a split brain scenario.
- **B. Kafka uses ZooKeeper's ephemeral node feature to elect a controller.**
- **C. The Controller uses the epoch number to prevent a split brain scenario.**
- **D. The Controller is responsible for electing Leaders among the partitions and replicas.**
- E. The Controller is responsible for reassigning partitions to the consumers in a Consumer Group.
- F. The number of Controllers should always be equal to the number of brokers alive in the cluster.

**Answer: B,C,D**

Explanation:

Kafka relies on ZooKeeper's ephemeral nodes to detect if a broker (controller) goes down and to elect a new controller.

The controller manages partition leadership assignments and handles leader election when a broker fails.

The epoch number ensures coordination and avoids outdated controllers acting on stale data.

### NEW QUESTION # 35

Your Kafka cluster has four brokers. The topic t1 on the cluster has two partitions, and it has a replication factor of three. You create a Consumer Group with four consumers, which subscribes to t1.

In the scenario above, how many Controllers are in the Kafka cluster?

- **A. One**
- B. two
- C. three
- D. Four

**Answer: A**

Explanation:

In a Kafka cluster, only one broker acts as the Controller at any given time. The Controller is responsible for managing cluster metadata, such as partition leadership and broker status. Even if the cluster has multiple brokers (in this case, four), only one is elected as the Controller, and others serve as regular brokers. If the current Controller fails, another broker is automatically elected to take its place.

