# Cloudera

*CDP-5001*
*CDP Administrator - Public Cloud Certification Exam*

**For More Information – Visit link below:**
**https://www.examsempire.com/**

**Product Version**

1. **Up to Date products, reliable and verified.**
2. **Questions and Answers in PDF Format.**


**https://examsempire.com/**

# Latest Version: 6.0

## Question: 1

You want to implement "predictive" autoscaling for a Data Service based on historical workload trends. What components and technologies would likely be involved in achieving this? (Choose two)

A. A time-series database to store and analyze historical metrics
B. Custom scripts integrated with Cloudera Manager's REST API for scaling actions
C. Apache Kafka for distributing scaling events to Data Hub worker nodes
D. Apache Ranger for role-based scaling control
E. HDFS audit logs for pattern analysis

**Answer: A,B**

## Question: 2

You're planning to use a queue depth metric from Apache Kafka as a trigger for autoscaling a Data Service. What considerations are important for implementing this effectively?

A. Ensure consistent Kafka consumer group naming for accurate metric aggregation.
B. Select a Data Service type designed for real-time message processing.
C. Thoroughly test scaling behavior to avoid oscillations due to rapid metric changes.
D. Configure Kerberos with short ticket lifetimes for security during scaling.
E. Disable HDFS replication for faster Data Service startup times.

**Answer: B,C**

## Question: 3

While autoscaling works well, you'd like to further optimize costs by making newly scaled instances "warm" - ready with cached data for faster processing. How could you achieve this? (Choose two)
A Integrate a distributed in-memory caching system with the Data Service.
B. Utilize HDFS transparent encryption to preload data with minimal overhead.
C. Increase the priority of the Data Service in YARN's 'fair' scheduler.
D. Employ a custom pre-scaling hook to pre-fetch data onto new worker nodes.
E. Configure Ozone storage with a higher replication factor during peak hours.

**Answer: A,D**

## Question: 4

You want to minimize disruption when autoscaling down a stateful Data Service. Which strategies would help? (Choose two)

A. Leverage a load balancer with session affinity for incoming requests.
B. Increase the autoscaling cooldown period to allow for graceful termination.
C. Design the Data Service to periodically persist its state to shared storage.
D. Deploy Apache Kafka as a message buffer in front of the Data Service.
E. Replicate the Data Service's data to HDFS for fault tolerance.

**Answer: B,C**

## Question: 5

A team requests highly dynamic autoscaling where the Data Service scales up and down within minutes. What factors should you evaluate to determine if this is feasible?

A. Virtual machine provisioning speed in your chosen cloud environment.
B. The Data Service's startup time and overhead for processing a single workload unit.
C. Network security group reconfiguration rules on the Data Hub cluster.
D. Apache Ranger policy granularity and refresh interval.
E. Cooldown periods in other Data Services on the same Data Hub.

**Answer: A,B**

## Question: 6

Your CDP environment hosts a Data Service that processes highly sensitive dat
a. Which security controls should be in place to mitigate the risk of data exposure? (Choose three)
A. Encrypt data at rest in HDFS or cloud object storage.
B. Apply strict access controls using Apache Ranger.
C. Regular security audits and penetration testing on the Data Service.
D. Enable verbose logging on the Data Service for anomaly detection.
E. Deploy the Data Service in an air-gapped network with no external connectivity.

**Answer: A,B,C**

## Question: 7

You're integrating a new Data Service with IDBroker for authentication and authorization. What are the key considerations for a secure setup? (Choose two)

A. Minimize the number of service principals and keytabs used by the Data Service.
B. Enforce short expiration times for Kerberos tickets.
C. Integrate Apache Knox as a gateway to the Data Service to consolidate security.
D. Grant the Data Service broad Ranger permissions for flexibility.
E. Replicate IDBroker data to multiple Data Hub clusters for redundancy.

**Answer: A,B**

## Question: 8

You need to enable fine-grained authorization on a Data Service that performs data transformations. Users should have varying levels of access to specific datasets and the ability to perform only certain actions. Which CDP component is best suited for this?

A. IDBroker
B. Apache Ranger
C. Apache Knox
D. Apache Atlas
E. Cloudera Manager

**Answer: B**

## Question: 9

A security review mandates the use of multi-factor authentication (MFA) for accessing Data Services within a CDP environment. Select the most common ways to integrate MF

A. Enforce password complexity and regular rotation within Ranger.
B. Configure hardware tokens for YARN container access.
C. Integrate IDBroker with an external identity provider that supports MFA.
D. Implement client-side certificate authentication for the Data Service.
E. Add custom MFA logic directly within the Data Service's code.

**Answer: C,D**

## Question: 10

You suspect a recently deployed Data Service has a vulnerability that could be exploited to gain access to other systems on the Data Hub cluster. What steps would you take to contain the issue and mitigate risk? (Choose two)

A. Immediately revoke Kerberos tickets issued to the Data Service.
B. Disable the Data Service in Cloudera Manager.
C. Isolate the Data Hub network using security groups or firewall rules.
D. Increase the replication factor of HDFS directories used by the Data Service.
E. Review container logs for the Data Service and look for anomalous activity.

**Answer: B,C**

# Thank You for Trying Our Product

**Special 16 USD Discount Coupon: NSZUBG3X**

**Email: support@examsempire.com**

**Check our Customer Testimonials and ratings available on every product page.**

**Visit our website.**

**https://examsempire.com/**