



**Free Questions for DAS-C01 by dumpsheet**

**Shared by Gay on 15-04-2024**

**For More Free Questions and Preparation Resources**

**Check the Links on Last Page**

# Question 1

---

## Question Type: MultipleChoice

---

A company is designing a data warehouse to support business intelligence reporting. Users will access the executive dashboard heavily each Monday and Friday morning

for 1 hour. These read-only queries will run on the active Amazon Redshift cluster, which runs on dc2.8xlarge compute nodes 24 hours a day, 7 days a week. There are

three queues set up in workload management: Dashboard, ETL, and System. The Amazon Redshift cluster needs to process the queries without wait time.

What is the MOST cost-effective way to ensure that the cluster processes these queries?

### Options:

---

- A-** Perform a classic resize to place the cluster in read-only mode while adding an additional node to the cluster.
- B-** Enable automatic workload management.
- C-** Perform an elastic resize to add an additional node to the cluster.
- D-** Enable concurrency scaling for the Dashboard workload queue.

### Answer:

---

D

## Question 2

---

**Question Type:** MultipleChoice

---

A company's data science team is designing a shared dataset repository on a Windows server. The data repository will store a large amount of training data that the data

science team commonly uses in its machine learning models. The data scientists create a random number of new datasets each day.

The company needs a solution that provides persistent, scalable file storage and high levels of throughput and IOPS. The solution also must be highly available and must

integrate with Active Directory for access control.

Which solution will meet these requirements with the LEAST development effort?

### Options:

---

**A-** Store datasets as files in an Amazon EMR cluster. Set the Active Directory domain for authentication.

**B-** Store datasets as files in Amazon FSx for Windows File Server. Set the Active Directory domain for authentication.

**C-** Store datasets as tables in a multi-node Amazon Redshift cluster. Set the Active Directory domain for authentication.

**D-** Store datasets as global tables in Amazon DynamoDB. Build an application to integrate authentication with the Active Directory domain.

**Answer:**

---

B

## Question 3

---

**Question Type: MultipleChoice**

---

A company is using an AWS Lambda function to run Amazon Athena queries against a cross-account AWS Glue Data Catalog. A query returns the following error:

HIVE METASTORE ERROR

The error message states that the response payload size exceeds the maximum allowed payload size. The queried table is already partitioned, and the data is stored in an

Amazon S3 bucket in the Apache Hive partition format.

Which solution will resolve this error?

### Options:

---

- A-** Modify the Lambda function to upload the query response payload as an object into the S3 bucket. Include an S3 object presigned URL as the payload in the Lambda function response.
- B-** Run the MSCK REPAIR TABLE command on the queried table.
- C-** Create a separate folder in the S3 bucket. Move the data files that need to be queried into that folder. Create an AWS Glue crawler that points to the folder instead of the S3 bucket.
- D-** Check the schema of the queried table for any characters that Athena does not support. Replace any unsupported characters with characters that Athena supports.

### Answer:

---

A

## Question 4

---

### Question Type: MultipleChoice

---

An IOT company is collecting data from multiple sensors and is streaming the data to Amazon Managed Streaming for Apache Kafka (Amazon MSK). Each sensor type has

its own topic, and each topic has the same number of partitions.

The company is planning to turn on more sensors. However, the company wants to evaluate which sensor types are producing the most data so that the company can scale

accordingly. The company needs to know which sensor types have the largest values for the following metrics: BytesInPerSec and MessagesInPerSec.

Which level of monitoring for Amazon MSK will meet these requirements?

**Options:**

---

- A- DEFAULT level
- B- PER TOPIC PER BROKER level
- C- PER BROKER level
- D- PER TOPIC level

**Answer:**

---

B

## Question 5

---

**Question Type: MultipleChoice**

---

An ecommerce company uses Amazon Aurora PostgreSQL to process and store live transactional data and uses Amazon Redshift for its data warehouse solution. A nightly ETL job has been implemented to update the Redshift cluster with new data from the PostgreSQL database. The business has grown rapidly and so has the size and cost of the Redshift cluster. The company's data analytics team needs to create a solution to archive historical data and only keep the most recent 12 months of data in Amazon

Redshift to reduce costs. Data analysts should also be able to run analytics queries that effectively combine data from live transactional data in PostgreSQL, current data in Redshift, and archived historical data.

Which combination of tasks will meet these requirements? (Select THREE.)

### Options:

---

- A-** Configure the Amazon Redshift Federated Query feature to query live transactional data in the PostgreSQL database.
- B-** Configure Amazon Redshift Spectrum to query live transactional data in the PostgreSQL database.
- C-** Schedule a monthly job to copy data older than 12 months to Amazon S3 by using the UNLOAD command, and then delete that data from the Redshift cluster. Configure Amazon Redshift Spectrum to access historical data in Amazon S3.
- D-** Schedule a monthly job to copy data older than 12 months to Amazon S3 Glacier Flexible Retrieval by using the UNLOAD command, and then delete that data from the Redshift cluster. Configure Redshift Spectrum to access historical data with S3 Glacier Flexible Retrieval.
- E-** Create a late-binding view in Amazon Redshift that combines live, current, and historical data from different sources.
- F-** Create a materialized view in Amazon Redshift that combines live, current, and historical data from different sources.

**Answer:**

---

A, C, E

## Question 6

---

**Question Type: MultipleChoice**

---

A company is creating a data lake by using AWS Lake Formation. The data that will be stored in the data lake contains sensitive customer information and must be encrypted at rest using an AWS Key Management Service (AWS KMS) customer managed key to meet regulatory requirements.

How can the company store the data in the data lake to meet these requirements?

**Options:**

---

- A-** Store the data in an encrypted Amazon Elastic Block Store (Amazon EBS) volume. Register the Amazon EBS volume with Lake Formation.
- B-** Store the data in an Amazon S3 bucket by using server-side encryption with AWS KMS (SSE-KMS). Register the S3 location with Lake Formation.
- C-** Encrypt the data on the client side and store the encrypted data in an Amazon S3 bucket. Register the S3 location with Lake Formation.



**D-** Store the data in an Amazon S3 Glacier Flexible Retrieval vault bucket. Register the S3 Glacier Flexible Retrieval vault with Lake Formation.

**Answer:**

---

B

## Question 7

---

**Question Type: MultipleChoice**

---

A large ecommerce company uses Amazon DynamoDB with provisioned read capacity and auto scaled write capacity to store its product catalog. The company uses Apache HiveQL statements on an Amazon EMR cluster to query the DynamoDB table. After the company announced a sale on all of its products, wait times for each query have increased. The data analyst has determined that the longer wait times are being caused by throttling when querying the table.

Which solution will solve this issue?

**Options:**

---

**A-** Increase the size of the EMR nodes that are provisioned.

- B-** Increase the number of EMR nodes that are in the cluster.
- C-** Increase the DynamoDB table's provisioned write throughput.
- D-** Increase the DynamoDB table's provisioned read throughput.

**Answer:**

---

D

## Question 8

---

**Question Type:** MultipleChoice

---

A company has a mobile app that has millions of users. The company wants to enhance the mobile app by including interactive data visualizations that show user trends.

The data for visualization is stored in a large data lake with 50 million rows. Data that is used in the visualization should be no more than two hours old.

Which solution will meet these requirements with the LEAST operational overhead?

**Options:**

---

- A-** Run an hourly batch process that renders user-specific data visualizations as static images that are stored in Amazon S3.
- B-** Precompute aggregated data hourly. Store the data in Amazon DynamoDB. Render the data by using the D3.js JavaScript library.
- C-** Embed an Amazon QuickSight Enterprise edition dashboard into the mobile app by using the QuickSight Embedding SDK. Refresh data in SPICE hourly.
- D-** Run Amazon Athena queries behind an Amazon API Gateway API. Render the data by using the D3.js JavaScript library.

**Answer:**

---

A

**To Get Premium Files for DAS-C01 Visit**

**<https://www.p2pexams.com/products/das-c01>**

**For More Free Questions Visit**

**<https://www.p2pexams.com/amazon/pdf/das-c01>**

