

AWS Certified Data Engineer - Associate (DEA-C01) exam training dumps & Data-Engineer-Associate valid test questions & AWS Certified Data Engineer - Associate (DEA-C01) test vce torrent



P.S. Free & New Data-Engineer-Associate dumps are available on Google Drive shared by DumpsFree:
https://drive.google.com/open?id=1rVkJOEKS-zsqd_uonIDFi5MUDO7IKW4X

The experts and professors of our company have designed the three different versions of the Data-Engineer-Associate study materials, including the PDF version, the online version and the software version. Now we are going to introduce the online version for you. There are a lot of advantages about the online version of the Data-Engineer-Associate Study Materials from our company. For instance, the online version can support any electronic equipment and it is not limited to all electronic equipment.

As one of the hot exam of our website, Amazon dumps pdf has a high pass rate which reach to 85%. According to our customer's feedback, our Data-Engineer-Associate vce braindumps covers mostly the same topics as included in the real exam. So if you practice our Data-Engineer-Associate Test Questions seriously and review test answers, pass exam will be absolute.

>> Data-Engineer-Associate Exam Simulator Fee <<

Data-Engineer-Associate Exam Material | Latest Data-Engineer-Associate Test Report

In order to serve you better, we have a complete system to you if you buy Data-Engineer-Associate study materials from us. We offer you free demo for you to have a try before buying. If you are satisfied with the exam, you can just add them to cart, and pay for it. You will obtain the downloading link and password for Data-Engineer-Associate Study Materials within ten minutes, if you don't, just contact us, we will solve the problem for you. After you buy, if you have some questions about the Data-Engineer-Associate exam braindumps after buying you can contact our service staff, they have the professional knowledge and will give you reply.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q69-Q74):

NEW QUESTION # 69

A company has a data warehouse in Amazon Redshift. To comply with security regulations, the company needs to log and store all user activities and connection activities for the data warehouse.

Which solution will meet these requirements?

- A. Create an Amazon S3 bucket. Enable logging for the Amazon Redshift cluster. Specify the S3 bucket in the logging configuration to store the logs.

- B. Create an Amazon Aurora MySQL database. Enable logging for the Amazon Redshift cluster. Write the logs to a table in the Aurora MySQL database.
- C. Create an Amazon Elastic Block Store (Amazon EBS) volume. Enable logging for the Amazon Redshift cluster. Write the logs to the EBS volume.
- D. Create an Amazon Elastic File System (Amazon EFS) file system. Enable logging for the Amazon Redshift cluster. Write logs to the EFS file system.

Answer: A

Explanation:

Problem Analysis:

The company must log all user activities and connection activities in Amazon Redshift for security compliance.

Key Considerations:

Redshift supports audit logging, which can be configured to write logs to an S3 bucket.

S3 provides durable, scalable, and cost-effective storage for logs.

Solution Analysis:

Option A: S3 for Logging

Standard approach for storing Redshift logs.

Easy to set up and manage with minimal cost.

Option B: Amazon EFS

EFS is unnecessary for this use case and less cost-efficient than S3.

Option C: Aurora MySQL

Using a database to store logs increases complexity and cost.

Option D: EBS Volume

EBS is not a scalable option for log storage compared to S3.

Final Recommendation:

Enable Redshift audit logging and specify an S3 bucket as the destination.

Reference:

[Amazon Redshift Audit Logging](#)

[Storing Logs in Amazon S3](#)

NEW QUESTION # 70

A retail company uses Amazon Aurora PostgreSQL to process and store live transactional data. The company uses an Amazon Redshift cluster for a data warehouse.

An extract, transform, and load (ETL) job runs every morning to update the Redshift cluster with new data from the PostgreSQL database. The company has grown rapidly and needs to cost optimize the Redshift cluster.

A data engineer needs to create a solution to archive historical data. The data engineer must be able to run analytics queries that effectively combine data from live transactional data in PostgreSQL, current data in Redshift, and archived historical data. The solution must keep only the most recent 15 months of data in Amazon Redshift to reduce costs.

Which combination of steps will meet these requirements? (Select TWO.)

- A. Schedule a monthly job to copy data that is older than 15 months to Amazon S3 by using the UNLOAD command. Delete the old data from the Redshift cluster. Configure Amazon Redshift Spectrum to access historical data in Amazon S3.
- B. Configure Amazon Redshift Spectrum to query live transactional data that is in the PostgreSQL database.
- C. Schedule a monthly job to copy data that is older than 15 months to Amazon S3 Glacier Flexible Retrieval by using the UNLOAD command. Delete the old data from the Redshift cluster. Configure Redshift Spectrum to access historical data from S3 Glacier Flexible Retrieval.
- D. Configure the Amazon Redshift Federated Query feature to query live transactional data that is in the PostgreSQL database.
- E. Create a materialized view in Amazon Redshift that combines live, current, and historical data from different sources.

Answer: A,D

Explanation:

The goal is to archive historical data from an Amazon Redshift data warehouse while combining live transactional data from Amazon Aurora PostgreSQL with current and historical data in a cost-efficient manner. The company wants to keep only the last 15 months of data in Redshift to reduce costs.

Option A: "Configure the Amazon Redshift Federated Query feature to query live transactional data that is in the PostgreSQL database." Redshift Federated Query allows querying live transactional data directly from Aurora PostgreSQL without having to move it into Redshift, thereby enabling seamless integration of the current data in Redshift and live data in PostgreSQL. This is a

cost-effective approach, as it avoids unnecessary data duplication.

Option C: "Schedule a monthly job to copy data that is older than 15 months to Amazon S3 by using the UNLOAD command. Delete the old data from the Redshift cluster. Configure Amazon Redshift Spectrum to access historical data in Amazon S3." This option uses Amazon Redshift Spectrum, which enables Redshift to query data directly in S3 without moving it into Redshift. By unloading older data (older than 15 months) to S3, and then using Spectrum to access it, this approach reduces storage costs significantly while still allowing the data to be queried when necessary.

Option B (Redshift Spectrum for live PostgreSQL data) is not applicable, as Redshift Spectrum is intended for querying data in Amazon S3, not live transactional data in Aurora.

Option D (S3 Glacier Flexible Retrieval) is not suitable because Glacier is designed for long-term archival storage with infrequent access, and querying data in Glacier for analytics purposes would incur higher retrieval times and costs.

Option E (materialized views) would not meet the need to archive data or combine it from multiple sources; it is best suited for combining frequently accessed data already in Redshift.

Reference:

[Amazon Redshift Federated Query](#)

[Amazon Redshift Spectrum Documentation](#)

[Amazon Redshift UNLOAD Command](#)

NEW QUESTION # 71

A company currently stores all of its data in Amazon S3 by using the S3 Standard storage class.

A data engineer examined data access patterns to identify trends. During the first 6 months, most data files are accessed several times each day. Between 6 months and 2 years, most data files are accessed once or twice each month. After 2 years, data files are accessed only once or twice each year.

The data engineer needs to use an S3 Lifecycle policy to develop new data storage rules. The new storage solution must continue to provide high availability.

Which solution will meet these requirements in the MOST cost-effective way?

- A. Transition objects to S3 Standard-Infrequent Access (S3 Standard-IA) after 6 months. Transfer objects to S3 Glacier Deep Archive after 2 years.
- B. Transition objects to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 6 months. Transfer objects to S3 Glacier Flexible Retrieval after 2 years.
- C. Transition objects to S3 Standard-Infrequent Access (S3 Standard-IA) after 6 months. Transfer objects to S3 Glacier Flexible Retrieval after 2 years.
- D. Transition objects to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 6 months. Transfer objects to S3 Glacier Deep Archive after 2 years.

Answer: A

Explanation:

To achieve the most cost-effective storage solution, the data engineer needs to use an S3 Lifecycle policy that transitions objects to lower-cost storage classes based on their access patterns, and deletes them when they are no longer needed. The storage classes should also provide high availability, which means they should be resilient to the loss of data in a single Availability Zone1. Therefore, the solution must include the following steps:

* Transition objects to S3 Standard-Infrequent Access (S3 Standard-IA) after 6 months. S3 Standard-IA is designed for data that is accessed less frequently, but requires rapid access when needed. It offers the same high durability, throughput, and low latency as S3 Standard, but with a lower storage cost and a retrieval fee2. Therefore, it is suitable for data files that are accessed once or twice each month. S3 Standard-IA also provides high availability, as it stores data redundantly across multiple Availability Zones1.

* Transfer objects to S3 Glacier Deep Archive after 2 years. S3 Glacier Deep Archive is the lowest-cost storage class that offers secure and durable storage for data that is rarely accessed and can tolerate a 12- hour retrieval time. It is ideal for long-term archiving and digital preservation3. Therefore, it is suitable for data files that are accessed only once or twice each year. S3 Glacier Deep Archive also provides high availability, as it stores data across at least three geographically dispersed Availability Zones1.

* Delete objects when they are no longer needed. The data engineer can specify an expiration action in the S3 Lifecycle policy to delete objects after a certain period of time. This will reduce the storage cost and comply with any data retention policies.

Option C is the only solution that includes all these steps. Therefore, option C is the correct answer.

Option A is incorrect because it transitions objects to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 6 months. S3 One Zone-IA is similar to S3 Standard-IA, but it stores data in a single Availability Zone. This means it has a lower availability and durability than S3 Standard-IA, and it is not resilient to the loss of data in a single Availability Zone1. Therefore, it does not provide high availability as required.

Option B is incorrect because it transfers objects to S3 Glacier Flexible Retrieval after 2 years. S3 Glacier Flexible Retrieval is a storage class that offers secure and durable storage for data that is accessed infrequently and can tolerate a retrieval time of minutes to hours. It is more expensive than S3 Glacier Deep Archive, and it is not suitable for data that is accessed only once or twice each

year3. Therefore, it is not the most cost-effective option.

Option D is incorrect because it combines the errors of option A and B. It transitions objects to S3 One Zone- IA after 6 months, which does not provide high availability, and it transfers objects to S3 Glacier Flexible Retrieval after 2 years, which is not the most cost-effective option.

:

- 1: Amazon S3 storage classes - Amazon Simple Storage Service
- 2: Amazon S3 Standard-Infrequent Access (S3 Standard-IA) - Amazon Simple Storage Service
- 3: Amazon S3 Glacier and S3 Glacier Deep Archive - Amazon Simple Storage Service
- [4]: Expiring objects - Amazon Simple Storage Service
- [5]: Managing your storage lifecycle - Amazon Simple Storage Service
- [6]: Examples of S3 Lifecycle configuration - Amazon Simple Storage Service
- [7]: Amazon S3 Lifecycle further optimizes storage cost savings with new features - What's New with AWS

NEW QUESTION # 72

A retail company uses an Amazon Redshift data warehouse and an Amazon S3 bucket. The company ingests retail order data into the S3 bucket every day.

The company stores all order data at a single path within the S3 bucket. The data has more than 100 columns. The company ingests the order data from a third-party application that generates more than 30 files in CSV format every day. Each CSV file is between 50 and 70 MB in size.

The company uses Amazon Redshift Spectrum to run queries that select sets of columns. Users aggregate metrics based on daily orders. Recently, users have reported that the performance of the queries has degraded. A data engineer must resolve the performance issues for the queries.

Which combination of steps will meet this requirement with LEAST developmental effort? (Select TWO.)

- A. Partition the order data in the S3 bucket based on order date.
- B. Configure the third-party application to create the files in JSON format.
- C. Configure the third-party application to create the files in a columnar format.
- D. Develop an AWS Glue ETL job to convert the multiple daily CSV files to one file for each day.
- E. Load the JSON data into the Amazon Redshift table in a SUPER type column.

Answer: A,C

Explanation:

The performance issue in Amazon Redshift Spectrum queries arises due to the nature of CSV files, which are row-based storage formats. Spectrum is more optimized for columnar formats, which significantly improve performance by reducing the amount of data scanned. Also, partitioning data based on relevant columns like order date can further reduce the amount of data scanned, as queries can focus only on the necessary partitions.

A . Configure the third-party application to create the files in a columnar format:

Columnar formats (like Parquet or ORC) store data in a way that is optimized for analytical queries because they allow queries to scan only the columns required, rather than scanning all columns in a row-based format like CSV.

Amazon Redshift Spectrum works much more efficiently with columnar formats, reducing the amount of data that needs to be scanned, which improves query performance.

Reference:

C . Partition the order data in the S3 bucket based on order date:

Partitioning the data on columns like order date allows Redshift Spectrum to skip scanning unnecessary partitions, leading to improved query performance.

By organizing data into partitions, you minimize the number of files Spectrum has to read, further optimizing performance.

Alternatives Considered:

B (Develop an AWS Glue ETL job): While consolidating files can improve performance by reducing the number of small files (which can be inefficient to process), it adds additional ETL complexity. Switching to a columnar format (Option A) and partitioning (Option C) provides more significant performance improvements with less development effort.

D and E (JSON-related options): Using JSON format or the SUPER type in Redshift introduces complexity and isn't as efficient as the proposed solutions, especially since JSON is not a columnar format.

Amazon Redshift Spectrum Documentation

Columnar Formats and Data Partitioning in S3

NEW QUESTION # 73

A company wants to analyze sales records that the company stores in a MySQL database. The company wants to correlate the

records with sales opportunities identified by Salesforce.

The company receives 2 GB of sales records every day. The company has 100 GB of identified sales opportunities. A data engineer needs to develop a process that will analyze and correlate sales records and sales opportunities. The process must run once each night.

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

- A. Use Amazon AppFlow to fetch sales opportunities from Salesforce. Use Amazon Kinesis Data Streams to fetch sales records from the MySQL database. Use Amazon Managed Service for Apache Flink to correlate the datasets. Use AWS Step Functions to orchestrate the process.
- B. Use Amazon AppFlow to fetch sales opportunities from Salesforce. Use AWS Glue to fetch sales records from the MySQL database. Correlate the sales records with sales opportunities. Use AWS Step Functions to orchestrate the process.
- C. Use Amazon Managed Workflows for Apache Airflow (Amazon MWAA) to fetch both datasets. Use AWS Lambda functions to correlate the datasets. Use AWS Step Functions to orchestrate the process.
- D. Use Amazon AppFlow to fetch sales opportunities from Salesforce. Use AWS Glue to fetch sales records from the MySQL database. Correlate the sales records with the sales opportunities. Use Amazon Managed Workflows for Apache Airflow (Amazon MWAA) to orchestrate the process.

Answer: B

Explanation:

Problem Analysis:

The company processes 2 GB of daily sales records and 100 GB of Salesforce sales opportunities.

The goal is to analyze and correlate the two datasets with low operational overhead.

The process must run once nightly.

Key Considerations:

Amazon AppFlow simplifies data integration with Salesforce.

AWS Glue can extract data from MySQL and perform ETL operations.

Step Functions can orchestrate workflows with minimal manual intervention.

Apache Airflow and Flink add complexity, which conflicts with the requirement for low operational overhead.

Solution Analysis:

Option A: MWAA + Lambda + Step Functions

Requires custom Lambda code for dataset correlation, increasing development and operational complexity.

Option B: AppFlow + Glue + MWAA

MWAA adds orchestration overhead compared to the simpler Step Functions.

Option C: AppFlow + Glue + Step Functions

AppFlow fetches Salesforce data, Glue extracts MySQL data, and Step Functions orchestrate the entire process.

Minimal setup and operational overhead, making it the best choice.

Option D: AppFlow + Kinesis + Flink + Step Functions

Using Kinesis and Flink for batch processing introduces unnecessary complexity.

Final Recommendation:

Use Amazon AppFlow to fetch Salesforce data, AWS Glue to process MySQL data, and Step Functions for orchestration.

Reference:

Amazon AppFlow Overview

AWS Glue ETL Documentation

AWS Step Functions

NEW QUESTION # 74

.....

Our Data-Engineer-Associate training materials are of high quality, and we also have free demo to help you know the content of the Data-Engineer-Associate exam dumps. Free update for 365 days after purchasing is available, and the update version will be sent to you timely. If you fail to pass the exam, we will return your money into the payment account. All we do is for your interest, and we also accept your suggestion and advice for Data-Engineer-Associate Training Materials.

Data-Engineer-Associate Exam Material: <https://www.dumpsfree.com/Data-Engineer-Associate-valid-exam.html>

Besides, Data-Engineer-Associate exam materials are compiled by skilled professionals, and they are familiar with the exam center, therefore the quality can be guaranteed. With our actual Amazon Data-Engineer-Associate questions PDF, Data-Engineer-Associate practice exams along with the support of our customer support team, you can be confident that you are getting the best possible Data-Engineer-Associate preparation material for the test. Most of the Data-Engineer-Associate study materials are written

by the famous experts in the field.

Working with Sound Titles and Credits, Performing Custom Validation: The 'CustomValidator' Control, Besides, Data-Engineer-Associate exam materials are compiled by skilled professionals, Data-Engineer-Associate and they are familiar with the exam center, therefore the quality can be guaranteed.

Maximize Your Chances of Getting Amazon Data-Engineer-Associate Exam Questions

With our actual Amazon Data-Engineer-Associate questions PDF, Data-Engineer-Associate practice exams along with the support of our customer support team, you can be confident that you are getting the best possible Data-Engineer-Associate preparation material for the test.

Most of the Data-Engineer-Associate study materials are written by the famous experts in the field, Data-Engineer-Associate simulating exam will inspire your potential, we will solve your problem by the first time and give you the most professional suggestions.

What's more, part of that DumpsFree Data-Engineer-Associate dumps now are free: https://drive.google.com/open?id=1rVkJAOEKS-zsqd_uonIDFi5MUDo7IKW4X