

100% Pass Quiz Data-Engineer-Associate - Pass-Sure Practice AWS Certified Data Engineer - Associate (DEA-C01) Exams



P.S. Free & New Data-Engineer-Associate dumps are available on Google Drive shared by Free4Torrent:
<https://drive.google.com/open?id=1InqtQbM7qJg3RhNaJi9EGDxOp1FZBZyF>

Nowadays, the certification has been one of the criteria for many companies to recruit employees. And in order to obtain the Data-Engineer-Associate certification, taking the Data-Engineer-Associate exam becomes essential. Although everyone hopes to pass the exam, the difficulties in preparing for it should not be overlooked. There are plenty of people who took a lot of energy and time but finally failed to pass. You really need our Data-Engineer-Associate practice materials which can work as the pass guarantee.

If you have never bought our Data-Engineer-Associate exam materials on the website before, we understand you may encounter many problems such as payment or downloading Data-Engineer-Associate practice quiz and so on, contact with us, we will be there. Our employees are diligent to deal with your need and willing to do their part on the Data-Engineer-Associate Study Materials. And they are trained specially and professionally to know every detail about our Data-Engineer-Associate learning prep.

>> Practice Data-Engineer-Associate Exams <<

Data-Engineer-Associate Latest Exam Review - Testking Data-Engineer-Associate Exam Questions

Our Data-Engineer-Associate study materials' developers to stand in the perspective of candidate, fully consider their material basis and actual levels of knowledge, formulated a series of scientific and reasonable learning mode, meet the conditions for each user to tailor their learning materials. What's more, our Data-Engineer-Associate Study Materials are cheap and cheap, and we buy more and deliver more. The more customers we buy, the bigger the discount will be. In order to make the user a better experience to the superiority of our Data-Engineer-Associate study materials.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q31-Q36):

NEW QUESTION # 31

A company uses Amazon Redshift as a data warehouse solution. One of the datasets that the company stores in Amazon Redshift contains data for a vendor.

Recently, the vendor asked the company to transfer the vendor's data into the vendor's Amazon S3 bucket once each week. Which solution will meet this requirement?

- A. Create an AWS Lambda function to connect to the Redshift data warehouse. Configure the Lambda function to use the Redshift COPY command to copy the required data to the vendor's S3 bucket on a schedule.
- B. Configure Amazon Redshift Spectrum to use the vendor's S3 bucket as destination. Enable dataquerying in both directions.

- C. Use the Amazon Redshift data sharing feature. Set the vendor's S3 bucket as the destination. Configure the source to be as a custom SQL query that selects the required data.
- D. Create an AWS Glue job to connect to the Redshift data warehouse. Configure the AWS Glue job to use the Redshift UNLOAD command to load the required data to the vendor's S3 bucket on a schedule.

Answer: D

Explanation:

The Redshift UNLOAD command is specifically designed to export query results to Amazon S3, and AWS Glue can orchestrate this as part of a scheduled job. This is the cleanest and most appropriate approach for recurring weekly data transfers:

"Use the Redshift UNLOAD command with AWS Glue to export data to Amazon S3. This pattern enables routine exports of selected data to external locations."

-Ace the AWS Certified Data Engineer - Associate Certification - version 2 - apple.pdf This avoids complexities of Redshift Spectrum or unsupported use of COPY commands in Lambda.

NEW QUESTION # 32

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. Create a materialized view in Amazon Redshift that combines live, current, and historical data from different sources.
- B. Configure the Amazon Redshift Federated Query feature 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 by using the UNLOAD command. Delete the old data from the Redshift cluster. Configure Amazon Redshift Spectrum to access historical data in Amazon S3.
- D. 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.
- E. Configure Amazon Redshift Spectrum to query live transactional data that is in the PostgreSQL database.

Answer: B,C

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.

References:

* Amazon Redshift Federated Query

* Amazon Redshift Spectrum Documentation

* Amazon Redshift UNLOAD Command

NEW QUESTION # 33

A company maintains a data warehouse in an on-premises Oracle database. The company wants to build a data lake on AWS. The company wants to load data warehouse tables into Amazon S3 and synchronize the tables with incremental data that arrives from the data warehouse every day.

Each table has a column that contains monotonically increasing values. The size of each table is less than 50 GB. The data warehouse tables are refreshed every night between 1 AM and 2 AM. A business intelligence team queries the tables between 10 AM and 8 PM every day.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Use an AWS Database Migration Service (AWS DMS) full load plus CDC job to load tables that contain monotonically increasing data columns from the on-premises data warehouse to Amazon S3. Use custom logic in AWS Glue to append the daily incremental data to a full-load copy that is in Amazon S3.
- B. Use an AWS Glue Java Database Connectivity (JDBC) connection. Configure a job bookmark for a column that contains monotonically increasing values. Write custom logic to append the daily incremental data to a full-load copy that is in Amazon S3.
- C. Use an AWS Database Migration Service (AWS DMS) full load migration to load the data warehouse tables into Amazon S3 every day Overwrite the previous day's full-load copy every day.
- D. Use AWS Glue to load a full copy of the data warehouse tables into Amazon S3 every day. Overwrite the previous day's full-load copy every day.

Answer: A

Explanation:

The company needs to load data warehouse tables into Amazon S3 and perform incremental synchronization with daily updates. The most efficient solution is to use AWS Database Migration Service (AWS DMS) with a combination of full load and change data capture (CDC) to handle the initial load and daily incremental updates.

Option A: Use an AWS Database Migration Service (AWS DMS) full load plus CDC job to load tables that contain monotonically increasing data columns from the on-premises data warehouse to Amazon S3. Use custom logic in AWS Glue to append the daily incremental data to a full-load copy that is in Amazon S3.

DMS is designed to migrate databases to AWS, and the combination of full load plus CDC is ideal for handling incremental data changes efficiently. AWS Glue can then be used to append the incremental data to the full data set in S3. This solution is highly operationally efficient because it automates both the full load and incremental updates.

Options B, C, and D are less operationally efficient because they either require writing custom logic to handle bookmarks manually or involve unnecessary daily full loads.

Reference:

[AWS Database Migration Service Documentation](#)

[AWS Glue Documentation](#)

NEW QUESTION # 34

A company stores logs in an Amazon S3 bucket. When a data engineer attempts to access several log files, the data engineer discovers that some files have been unintentionally deleted.

The data engineer needs a solution that will prevent unintentional file deletion in the future.

Which solution will meet this requirement with the LEAST operational overhead?

- A. Configure replication for the S3 bucket.
- B. Use an Amazon S3 Glacier storage class to archive the data that is in the S3 bucket.
- C. Manually back up the S3 bucket on a regular basis.
- D. **Enable S3 Versioning for the S3 bucket.**

Answer: D

Explanation:

To prevent unintentional file deletions and meet the requirement with minimal operational overhead, enabling S3 Versioning is the best solution.

* S3 Versioning:

* S3 Versioning allows multiple versions of an object to be stored in the same S3 bucket. When a file is deleted or overwritten, S3 preserves the previous versions, which means you can recover from accidental deletions or modifications.

* Enabling versioning requires minimal overhead, as it is a bucket-level setting and does not require additional backup processes or data replication.

* Users can recover specific versions of files that were unintentionally deleted, meeting the needs of the data engineer to avoid accidental data loss.

NEW QUESTION # 35

An airline company is collecting metrics about flight activities for analytics. The company is conducting a proof of concept (POC) test to show how analytics can provide insights that the company can use to increase on-time departures.

The POC test uses objects in Amazon S3 that contain the metrics in .csv format. The POC test uses Amazon Athena to query the data. The data is partitioned in the S3 bucket by date.

As the amount of data increases, the company wants to optimize the storage solution to improve query performance.

Which combination of solutions will meet these requirements? (Choose two.)

- A. Add a randomized string to the beginning of the keys in Amazon S3 to get more throughput across partitions.
- B. Use an S3 bucket that is in the same AWS Region where the company runs Athena queries.
- C. Use an S3 bucket that is in the same account that uses Athena to query the data.
- D. Preprocess the .csv data to JSON format by fetching only the document keys that the query requires.
- E. Preprocess the .csv data to Apache Parquet format by fetching only the data blocks that are needed for predicates.

Answer: B,E

Explanation:

Using an S3 bucket that is in the same AWS Region where the company runs Athena queries can improve query performance by reducing data transfer latency and costs. Preprocessing the .csv data to Apache Parquet format can also improve query performance by enabling columnar storage, compression, and partitioning, which can reduce the amount of data scanned and fetched by the query. These solutions can optimize the storage solution for the POC test without requiring much effort or changes to the existing data pipeline. The other solutions are not optimal or relevant for this requirement. Adding a randomized string to the beginning of the keys in Amazon S3 can improve the throughput across partitions, but it can also make the data harder to query and manage. Using an S3 bucket that is in the same account that uses Athena to query the data does not have any significant impact on query performance, as long as the proper permissions are granted. Preprocessing the .csv data to JSON format does not offer any benefits over the .csv format, as both are row-based and verbose formats that require more data scanning and fetching than columnar formats like Parquet. Reference:

Best Practices When Using Athena with AWS Glue

Optimizing Amazon S3 Performance

AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide

NEW QUESTION # 36

.....

As you know the registration fee for the AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) certification exam is itself very high, varying between \$100 and \$1000. And after paying the registration fee for better preparation a candidate needs budget-friendly and reliable AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) pdf questions. That is why Free4Torrent has compiled the most reliable updated Data-Engineer-Associate Exam Questions with up to 1 year of free updates. The Amazon Data-Engineer-Associate practice test can be used right after being bought by the customer and they can avail of the benefits given in the AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) pdf questions.

Data-Engineer-Associate Latest Exam Review: <https://www.free4torrent.com/Data-Engineer-Associate-braindumps-torrent.html>

Usually the recommended Data-Engineer-Associate Latest Exam Review - AWS Certified Data Engineer - Associate (DEA-C01) dumps demo get you bored and you lose interest in irrelevant lengthy details, And we are just right here to give you help to pass the Data-Engineer-Associate exam, Amazon Practice Data-Engineer-Associate Exams We just provide the actual test latest version and key questions, The specialized knowledge with Data-Engineer-Associate exam prep files are your foundation of foothold in this competitive society.

Stuart's purchase of his new milk production Data-Engineer-Associate Latest Exam Review process were hardly auspicious, Managing Technical Debt: Reducing Friction in Software Development, Usually the recommended AWS Certified Data Engineer - Associate (DEA-C01) Data-Engineer-Associate Dumps demo get you bored and you lose interest in irrelevant lengthy details.

100% Pass Quiz 2026 Amazon Unparalleled Practice Data-Engineer-

Associate Exams

And we are just right here to give you help to pass the Data-Engineer-Associate exam. We just provide the actual test latest version and key questions, The specialized knowledge with Data-Engineer-Associate exam prep files are your foundation of foothold in this competitive society.

These Amazon Data-Engineer-Associate exam question formats are PDF dumps, web-based practice test software, and web-based practice test software.

P.S. Free & New Data-Engineer-Associate dumps are available on Google Drive shared by Free4Torrent:

<https://drive.google.com/open?id=1InqtQbM7qJg3RhNaJi9EGDxOp1FZBZyF>