

Brain DEA-C01 Exam Will Be Your Wisest Choice to Pass SnowPro Advanced: Data Engineer Certification Exam



BTW, DOWNLOAD part of PassExamDumps DEA-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=1ahFhImx4sVIQJODnT42vGytaTxaql0BZ>

With DEA-C01 study materials, you will have more flexible learning time. With DEA-C01 study materials, you can flexibly arrange your study time according to your own life. You don't need to be in a hurry to go to classes after work as the students who take part in a face-to-face class, and you also never have to disrupt your schedule for learning. DEA-C01 Study Materials help you not only to avoid all the troubles of learning but also to provide you with higher learning quality than other students'.

Snowflake DEA-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	• Security: The Security topic of the DEA-C01 test covers the principles of Snowflake security, including the management of system roles and data governance. It measures the ability to secure data and ensure compliance with policies, crucial for maintaining secure data environments for Snowflake Data Engineers and Software Engineers.
Topic 2	• Data Transformation: The SnowPro Advanced: Data Engineer exam evaluates skills in using User-Defined Functions (UDFs), external functions, and stored procedures. It assesses the ability to handle semi-structured data and utilize Snowpark for transformations. This section ensures Snowflake engineers can effectively transform data within Snowflake environments, critical for data manipulation tasks.
Topic 3	• Storage and Data Protection: The topic tests the implementation of data recovery features and the understanding of Snowflake's Time Travel and micro-partitions. Engineers are evaluated on their ability to create new environments through cloning and ensure data protection, highlighting essential skills for maintaining Snowflake data integrity and accessibility.
Topic 4	• Data Movement: Snowflake Data Engineers and Software Engineers are assessed on their proficiency to load, ingest, and troubleshoot data in Snowflake. It evaluates skills in building continuous data pipelines, configuring connectors, and designing data sharing solutions.
Topic 5	• Performance Optimization: This topic assesses the ability to optimize and troubleshoot underperforming queries in Snowflake. Candidates must demonstrate knowledge in configuring optimal solutions, utilizing caching, and monitoring data pipelines. It focuses on ensuring engineers can enhance performance based on specific scenarios, crucial for Snowflake Data Engineers and Software Engineers.

DEA-C01 Valid Braindumps Questions, DEA-C01 Latest Test Vce

We have three versions of our DEA-C01 exam questions: the PDF, Software and APP online. Because our PDF version of the learning material is available for customers to print, so that your free time is fully utilized. Everything you do will help you pass the DEA-C01 Exam and get your DEA-C01 certificate. Of course, the APP and PC versions are also very popular. They can simulate the actual operation of the test environment, and users can perform mock tests for a limited time.

Snowflake SnowPro Advanced: Data Engineer Certification Exam Sample Questions (Q297-Q302):

NEW QUESTION # 297

A data engineer has a one-time task to read data from objects that are in Apache Parquet format in an Amazon S3 bucket. The data engineer needs to query only one column of the data.

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

- A. Configure an AWS Lambda function to load data from the S3 bucket into a pandas dataframe. Write a SQL SELECT statement on the dataframe to query the required column.
- B. Prepare an AWS Glue DataBrew project to consume the S3 objects and to query the required column.
- **C. Use S3 Select to write a SQL SELECT statement to retrieve the required column from the S3 objects.**
- D. Run an AWS Glue crawler on the S3 objects. Use a SQL SELECT statement in Amazon Athena to query the required column.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/storage-inventory-athena-query.html> S3 Select allows you to retrieve a subset of data from an object stored in S3 using simple SQL expressions. It is capable of working directly with objects in Parquet format.

NEW QUESTION # 298

A company uses Amazon Redshift to store order transactions from the current day. The company has an orders table that contains the previous order data. The company also has a staging table that contains new or updated order records.

The company needs to remove stale records from the orders table and insert the most recent data in the orders table from the staging table. Several downstream applications need the orders table to display up-to-date information.

Which solution will meet these requirements?

- A. Unload the orders table and the staging table to Amazon S3. Delete stale orders table data and insert new staging table data in Amazon S3 by using Amazon Athena. Copy the orders S3 table to the orders Amazon Redshift table.
- **B. Write an Amazon Redshift stored procedure that deletes the stale records from the orders table and inserts new records from the staging table.**
- C. Use Amazon Redshift Spectrum to delete stale records from the orders table and insert records from the staging table into the orders table.
- D. Use Amazon Athena federated queries to read stale records from the orders table. Delete the stale records and insert the records from the staging table into the orders table.

Answer: B

Explanation:

An Amazon Redshift stored procedure can perform the delete-and-insert logic directly inside the warehouse in a single controlled operation. This is the most appropriate way to remove stale rows from the target table and load the latest records from the staging table while keeping the orders table current for downstream applications, without introducing unnecessary data movement or external processing.

NEW QUESTION # 299

Does sensitive data in Snowflake is modified in an existing table while applying Masking policies?

- A. NO
- B. YES

Answer: A

Explanation:

Explanation

Snowflake supports masking policies as a schema-level object to protect sensitive data from unauthorized access while allowing authorized users to access sensitive data at query runtime. This means that sensitive data in Snowflake is not modified in an existing table (i.e. no static masking). Rather, when users execute a query in which a masking policy applies, the masking policy conditions determine whether unauthorized users see masked, partially masked, obfuscated, or tokenized data.

NEW QUESTION # 300

A retail company stores customer data in an Amazon S3 bucket. Some of the customer data contains personally identifiable information (PII) about customers. The company must not share PII data with business partners.

A data engineer must determine whether a dataset contains PII before making objects in the dataset available to business partners. Which solution will meet this requirement with the LEAST manual intervention?

- A. Create an AWS Lambda function to identify PII in S3 objects. Schedule the function to run periodically.
- B. Create a table in AWS Glue Data Catalog. Write custom SQL queries to identify PII in the table. Use Amazon Athena to run the queries.
- C. Configure AWS CloudTrail to monitor S3 PUT operations. Inspect the CloudTrail trails to identify operations that save PII.
- D. Configure the S3 bucket and S3 objects to allow access to Amazon Macie. Use automated sensitive data discovery in Macie.

Answer: D

Explanation:

Amazon Macie is a fully managed data security service that automatically discovers, classifies, and protects sensitive data such as personally identifiable information (PII) in Amazon S3.

Macie's automated sensitive data discovery feature can scan S3 objects for PII without requiring manual intervention, making it the most efficient and automated solution for identifying sensitive data before sharing it with business partners.

CloudTrail is useful for monitoring API calls but is not designed to inspect the contents of S3 objects for PII. This would require manual effort to inspect the data saved in S3, which is inefficient.

While Lambda could be used to scan for PII, it would require custom code to identify sensitive data. This approach increases manual effort for development and maintenance compared to using a fully managed service like Macie.

Writing custom SQL queries to identify PII is a manual process and requires the creation and management of Glue tables and Athena queries. This approach adds complexity and manual intervention compared to the automated capabilities of Macie.

NEW QUESTION # 301

Which output is provided by both the `SYSTEM$CLUSTERING_DEPTH` function and the `SYSTEM$CLUSTERING_INFORMATION` function?

- A. `average_depth`
- B. `average_overlaps`
- C. `total_partition_count`
- D. `notes`

Answer: A

Explanation:

Explanation

The output that is provided by both the `SYSTEM$CLUSTERING_DEPTH` function and the `SYSTEM$CLUSTERING_INFORMATION` function is `average_depth`. This output indicates the average number of micro-partitions that contain data for a given column value or combination of column values. The other outputs are not common to both functions. The `notes` output is only provided by the `SYSTEM$CLUSTERING_INFORMATION` function and it contains additional information or recommendations about the clustering status of the table. The `average_overlaps` output is only provided by the `SYSTEM$CLUSTERING_DEPTH` function and it indicates the average number of micro-partitions that overlap with other micro-

partitions for a given column value or combination of column values. The `total_partition_count` output is only provided by the `SYSTEM$CLUSTERING_INFORMATION` function and it indicates the total number of micro-partitions in the table.

• • • • •

- [illegible]