Hot Valid DEA-C02 Test Camp Supply you Free-Download Study Test for DEA-C02: SnowPro Advanced: Data Engineer (DEA-C02) to Study casually

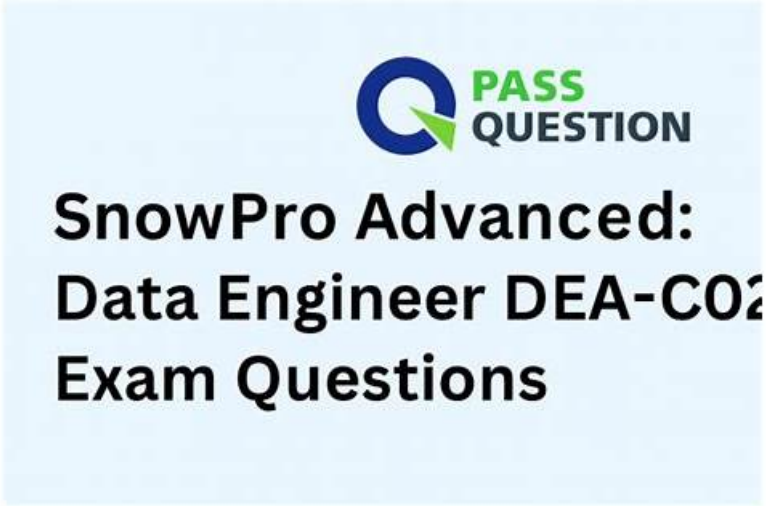

BONUS!!! Download part of PracticeMaterial DEA-C02 dumps for free: https://drive.google.com/open?id=1IBjJy3lZ3RZpoGFU9BNo9_k5iQ2Z06wH

You only need 20-30 hours to practice our software materials and then you can attend the exam. It costs you little time and energy. The DEA-C02 exam questions are easy to be mastered and simplified the content of important information. The DEA-C02 test guide conveys more important information with amount of answers and questions, thus the learning for the examinee is easy and highly efficient. So it is convenient for the learners to master the DEA-C02 Guide Torrent and pass the DEA-C02 exam in a short time.

It is well known that even the best people fail sometimes, not to mention the ordinary people. In face of the DEA-C02 exam, everyone stands on the same starting line, and those who are not excellent enough must do more. Every year there are a large number of people who can't pass the DEA-C02 Exam smoothly. But we are professional in this career for over ten years. And our DEA-C02 study materials will help you pass the exam easily.

>> Valid DEA-C02 Test Camp <<

PracticeMaterial Make its Snowflake DEA-C02 Exam Questions Engaging

We understand our candidates have no time to waste, everyone wants an efficient learning. So we take this factor into consideration, develop the most efficient way for you to prepare for the DEA-C02 exam, that is the real questions and answers practice mode, firstly, it simulates the real DEA-C02 test environment perfectly, which offers greatly help to our customers. Secondly, it includes printable PDF Format of DEA-C02 Exam Questions, also the instant access to download make sure you can study anywhere and anytime. All in all, high efficiency of DEA-C02 exam material is the reason for your selection.

Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Sample Questions (Q352-Q357):

NEW QUESTION # 352
You are designing a complex data pipeline in Snowflake that involves multiple interdependent Tasks. Several of these Tasks need to access sensitive customer data, and you want to ensure that the least privilege principle is followed. How should you configure the Tasks and their associated roles to minimize the risk of unauthorized data access while maintaining the functionality of the pipeline? (Select TWO)

- A. Use stored procedures executed with 'EXECUTE AS CALLER to encapsulate the sensitive data access logic. The stored procedure owner (who should have appropriate privileges) grants execute privilege to the Task's role, but the Task itself does

not directly interact with the sensitive data.

- B. Create separate, specific roles for each Task or group of related Tasks. Grant each role only the minimum necessary privileges to access the specific tables and functions required by that Task. Assign each Task to the appropriate role using the 'EXECUTE AS OWNER clause.o
- C. Grant the role to the user that owns all the Tasks. This ensures that all Tasks have the necessary privileges to access any data within the Snowflake account.
- D. Create a single role with broad data access privileges and grant this role to all Tasks. This simplifies role management and ensures that no Task encounters permission errors during execution.
- E. Grant 'SELECT privilege on all tables containing sensitive data to the 'PUBLIC' role. Tasks will inherit these privileges and can access the data without explicit role assignments.

Answer: A,B

Explanation:
Following the least privilege principle is crucial for security. Creating specific roles for each Task (B) and granting only the necessary privileges minimizes the risk of unauthorized access. Using stored procedures with 'EXECUTE AS CALLER (C) allows you to control data access through a separate, secured object, further limiting the Task's direct access to sensitive data. Option A grants excessive privileges, violating least privilege. Option D simplifies management but compromises security. Option E makes sensitive data broadly accessible, which is highly insecure.

NEW QUESTION # 353
A data engineering team is building a real-time dashboard in Snowflake to monitor website traffic. The dashboard relies on a complex query that joins several large tables. The query execution time is consistently exceeding the acceptable threshold, impacting dashboard responsiveness. Historical data is stored in a separate table and rarely changes. You suspect caching is not being utilized effectively. Which of the following actions would BEST improve the performance of this dashboard and leverage Snowflake's caching features?

- A. Materialize the historical data into a separate table that utilizes clustering and indexing for faster query performance. Refresh this table periodically.
- B. Use 'RESULT_SCAN' to cache the query result in the user session for subsequent queries. This is especially effective for large datasets that don't change frequently.
- C. Increase the size of the virtual warehouse. A larger warehouse will have more resources to execute the query, and the results will be cached for a longer period.
- D. Replace the complex query with a series of simpler queries. This will reduce the amount of data that needs to be processed at any one time.
- E. Create a materialized view that pre-computes the results of the complex query. Snowflake will automatically refresh the materialized view when the underlying data changes.

Answer: E

Explanation:
Materialized views are the best option in this scenario. They pre-compute the results of the complex query and store them in a separate table. Snowflake automatically refreshes the materialized view when the underlying data changes, ensuring that the dashboard always displays the most up-to-date information. While increasing the virtual warehouse size (D) can help initially, it's a more expensive and less targeted solution. 'RESULT_SCAN' (A) is session-specific and not suitable for persistent caching for a dashboard accessed by multiple users. Materializing the historical data (B) might help, but it doesn't address the core issue of the complex query. Breaking the query into smaller parts (E) might not be efficient and can introduce complexity.

NEW QUESTION # 354
A financial institution is using Snowflake to store transaction data for millions of customers. The data is stored in a table named 'TRANSACTIONS with columns such as 'TRANSACTION ID, 'CUSTOMER ID', 'TRANSACTION DATE, 'TRANSACTION_AMOUNT, and 'MERCHANT CATEGORY'. Analysts are running complex analytical queries that often involve filtering transactions by 'TRANSACTION_DATE, 'MERCHANT CATEGORY' , and 'TRANSACTION_AMOUNT ranges. These queries are experiencing performance bottlenecks. The data team wants to leverage query acceleration service to improve performance without significantly altering the existing query patterns. Which of the following actions or combination of actions would be MOST beneficial, considering the constraints and the nature of the queries? (Select TWO)

- A. Create materialized views pre-aggregating the transaction data by 'MERCHANT_CATEGORY and 'TRANSACTION_DATE, and enable query acceleration on the virtual warehouse.

- B. Create separate virtual warehouses dedicated to reporting queries and ad-hoc queries respectively. Enable query acceleration only for the warehouse running reporting queries.
- C. Enable Search Optimization Service for the 'TRANSACTIONS' table, specifically targeting the 'MERCHANT_CATEGORY column. Enable query acceleration on the virtual warehouse.
- D. Enable Automatic Clustering on the 'TRANSACTIONS' table, ordering the keys as 'TRANSACTION_DATE, 'MERCHANT_CATEGORY', 'CUSTOMER_ID. Then, enable query acceleration on the virtual warehouse.
- E. Increase the size of the virtual warehouse used for running the queries and enable query acceleration on the warehouse without further modifications.

Answer: C,D

Explanation:
Enabling Automatic Clustering on 'TRANSACTIONS with the specified key order ('TRANSACTION DATES, 'MERCHANT_CATEGORY , 'CUSTOMER_ID') aligns the data layout with common query patterns, allowing Snowflake to efficiently prune irrelevant data during query execution. This drastically improves query performance. Enabling Search Optimization on the 'MERCHANT_CATEGORY further enhances query performance by creating search access paths that enable faster lookups and filtering based on merchant category. Simply increasing the warehouse size (option A) may provide some improvement, but it's less targeted and potentially less cost-effective than optimizing the data organization. While dedicated warehouses (option C) can improve concurrency, they do not address the underlying performance bottleneck related to data access. Materialized views (option E) can be beneficial, but they require careful design and maintenance, and they might not be flexible enough for ad-hoc queries with varying filter conditions. Clustering and search optimization provide a more general and efficient solution in this scenario.

NEW QUESTION # 355
You have a Snowflake Stream named 'ORDERS STREAM' on an 'ORDERS' table, which is used to incrementally load data into a historical orders table named 'HISTORICAL ORDERS'. The data pipeline involves a series of tasks: 1) Consume changes from the 'ORDERS STREAM', 2) Apply transformations and data quality checks, and 3) Merge the changes into 'HISTORICAL ORDERS' using a MERGE statement. After a recent data load, you notice that the 'HISTORICAL ORDERS' table contains duplicate records for certain 'ORDER values. The MERGE statement uses 'ORDER ID' as the matching key. You have confirmed that the transformation logic is correct and idempotent. Examine the MERGE statement below. What could be causing the duplicates, given the context of Streams and incremental loading?

- A. The 'ORDERS STREAM' is retaining historical data beyond the data retention period, causing older records to be re-processed.
- B. Multiple tasks are concurrently consuming from the same 'ORDERS STREAM' without proper coordination, causing records to be processed multiple times.
- C. The stream is not configured to capture DELETE operations from the ORDERS table, causing records that should have been removed in HISTORICAL ORDERS to remain.
- D. The MERGE statement is not correctly handling updates and deletes from the stream. The 'WHEN NOT MATCHED' and 'WHEN MATCHED' clauses are not mutually exclusive, leading to potential insertions of duplicate rows.
- E. The stream's or 'BEFORE clause is being used incorrectly, potentially rewinding the stream to an earlier point in time.

Answer: B

Explanation:
The most likely cause of duplicate records, given the correct transformation logic and idempotent behavior, is D (Concurrent consumption from the stream) . If multiple tasks or processes are consuming from the same stream without proper coordination , they can both read the same changes and apply them to the 'HISTORICAL ORDERS' table. Option A (Retention Period) would cause data loss , not duplication, because older changes are simply lost. Option B is unlikely if you're not explicitly using or 'BEFORE' and even if you were, a simple stream consumption would advance the offset. If that clause was incorrectly configured, it would still not cause duplicates. Option C is possible but less likely if the transformation logic has been verified. If so many records are malformed that MERGE is broken, more evidence than duplication would be expected. Duplication alone with idempotent transformation suggests consumption error. Option E is a factor that could lead to 'HISTORICAL ORDERS' diverging from 'ORDERS but not a cause duplication.

NEW QUESTION # 356
You are designing a data loading process for a high-volume streaming data source. The data arrives as Avro files in an AWS S3 bucket. You need to load this data into a Snowflake table with minimal latency and operational overhead. Which of the following combinations of Snowflake features and configurations would be MOST suitable for this scenario? (Select TWO)

- A. Implement Snowpipe with auto-ingest configured to listen for S3 event notifications whenever a new Avro file is added to the bucket.
- B. Create a custom Spark application that reads Avro files from S3, transforms the data, and then writes it to Snowflake using the Snowflake Spark connector.
- C. Use a Kafka connector to stream data directly from the Kafka topic to Snowflake.
- D. Configure an external table pointing to the S3 bucket and query the Avro files directly from Snowflake.
- E. Use the 'COPY INTO' command with a scheduled task that runs every 5 minutes to load new files from the S3 bucket.

Answer: A,C

Explanation:
Explanation: Options B and C offer the best combination of low latency and operational efficiency for streaming data. Snowpipe with auto-ingest provides near real-time loading triggered by S3 events. A Kafka connector provides a direct data stream to Snowflake. Option A introduces latency due to the scheduled task interval and doesn't scale well for high-volume streams. Option D adds operational overhead with Spark application management. Option E is suitable for ad-hoc querying but not ideal for continuous data loading.

NEW QUESTION # 357
.....

We offer free demo DEA-C02 questions answers and trial services at PracticeMaterial. You can always check out our DEA-C02 certification exam dumps questions that will help you pass the DEA-C02 exams. With our well-researched and well-curated exam DEA-C02 dumps, you can surely pass the exam in the best marks. We continuously update our products by adding latest questions in our DEA-C02 Pdf Files. After the date of purchase, you will receive free updates for one year. You will also be able to get discounts for DEA-C02 on complete packages.

Study DEA-C02 Test: https://www.practicematerial.com/DEA-C02-exam-materials.html

For some difficult points of the DEA-C02 exam questions which you may feel hard to understand or easy to confuse for too similar with the others, After downloading and installing, Soft version of DEA-C02 VCE dumps can be used and copied into other computer offline, Every one customer who uses our Snowflake DEA-C02 test cram has no worries about the passing for the goal of each staff in our company is making the candidates get though the DEA-C02 test with 100% certainty, Snowflake Valid DEA-C02 Test Camp Then you can do a detail study plan and the success will be a little case.

Set its Swipe attribute to Up, For the purpose DEA-C02 of this book, we want to keep our graph simple, For some difficult pointsof the DEA-C02 Exam Questions which you may feel hard to understand or easy to confuse for too similar with the others.

Find Success In Exam With Snowflake DEA-C02 PDF Questions

After downloading and installing, Soft version of DEA-C02 VCE dumps can be used and copied into other computer offline, Every one customer who uses our Snowflake DEA-C02 test cram has no worries about the passing for the goal of each staff in our company is making the candidates get though the DEA-C02 test with 100% certainty.

Then you can do a detail study plan and the success will be a little case, We have the leading brand in this carrer and successfully help tens of thousands of our customers pass therir DEA-C02 exam and get admired certification.

Materials

- 100% Pass 2026 Marvelous Snowflake Valid DEA-C02 Test Camp □ 《 www.pdfvce.com 》 is best website to obtain ▷ DEA-C02 ◁ for free download □Dump DEA-C02 Check
- Reliable DEA-C02 Braindumps Pdf □ Valid Braindumps DEA-C02 Questions □ Simulation DEA-C02 Questions ➽ Search for [DEA-C02] and download exam materials for free through ✔ www.exam4labs.com □✔□ □Test DEA-C02 Online
- Polish Your Abilities To Easily Get the Snowflake DEA-C02 Certification □ Download { DEA-C02 } for free by simply entering " www.pdfvce.com " website □Dump DEA-C02 Check
- Free PDF Snowflake - DEA-C02 –Reliable Valid Test Camp □ Open ➥ www.examcollectionpass.com □ and search for ➤ DEA-C02 □ to download exam materials for free □DEA-C02 Exam Questions Answers
- myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.divephotoguide.com, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, Disposable vapes

P.S. Free 2026 Snowflake DEA-C02 dumps are available on Google Drive shared by PracticeMaterial: https://drive.google.com/open?id=1IBjJy3lZ3RZpoGFU9BNo9_k5iQ2Z06wH