

Study Snowflake DEA-C02 Group - DEA-C02 Latest Exam Fee



P.S. Free & New DEA-C02 dumps are available on Google Drive shared by VCEEngine: <https://drive.google.com/open?id=1o1aZQecCmpvgUbyLQEY56LSRojXLFEG>

The development of science and technology makes our life more comfortable and convenient, which also brings us more challenges. Many company requests candidates not only have work experiences, but also some professional certifications. Therefore it is necessary to get a professional Snowflake certification to pave the way for a better future. The DEA-C02 question and answers produced by our company, is helpful for our customers to pass their DEA-C02 exams and get the DEA-C02 certification within several days. Our DEA-C02 exam questions are your best choice.

In order to meet the different demands of the different customers, these experts from our company have designed three different versions of the DEA-C02 reference guide. All customers have the right to choose the most suitable version according to their need. The PDF version of the DEA-C02 exam prep has many special functions, including download the demo for free, support the printable format and so on. We can make sure that the PDF version of the DEA-C02 Test Questions will be very convenient for all people. Of course, if you choose our DEA-C02 study materials, you will love it.

>> Study Snowflake DEA-C02 Group <<

Most probable real and updated Snowflake DEA-C02 exam questions

Our DEA-C02 test questions provide free trial services for all customers so that you can better understand our products. You can experience the effects of outside products in advance by downloading clue versions of our DEA-C02 exam torrent. In addition, it has simple procedure to buy our learning materials. After your payment is successful, you will receive an e-mail from our company within 10 minutes. In a matter of seconds, you will receive an assessment report based on each question you have practiced on our DEA-C02 test material. The final result will show you the correct and wrong answers so that you can understand your learning ability so that you can arrange the learning tasks properly and focus on the targeted learning tasks with DEA-C02 test questions. So you can understand the wrong places and deepen the impression of them to avoid making the same mistake again.

Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Sample Questions (Q60-Q65):

NEW QUESTION # 60

A data engineering team is managing a Snowflake warehouse that supports a high volume of ad-hoc queries from data analysts exploring a large, semi-structured JSON dataset containing website clickstream data¹. The query performance is frequently slow, and analysts are complaining about long wait times. The warehouse is already sized appropriately. You have identified that many of the queries filter on nested JSON attributes that are not explicitly indexed. Considering only query acceleration service features, what is the MOST effective approach to improve query performance for these ad-hoc queries without modifying the queries themselves or significantly increasing storage costs?

- A. Use a combination of materialized views and query acceleration for best performance.
- B. **Enable Search Optimization Service for the table containing the JSON data. Selectively enable search optimization on the specific columns and nested paths that are frequently used in WHERE clause predicates.**
- C. Enable the Materialized View feature to create materialized views over the frequently queried nested JSON attributes.
- D. Create a dedicated virtual warehouse specifically for ad-hoc queries, and enable query acceleration on this warehouse.
- E. Enable Automatic Clustering on the table to improve data organization based on query patterns.

Answer: B

Explanation:

Search Optimization Service is designed specifically to improve query performance on semi-structured data and complex predicates, especially on JSON data. It automatically creates and maintains search access paths, including paths for nested JSON attributes, enabling faster filtering and retrieval of relevant data. Materialized Views are beneficial, but require creation and maintenance, and might not be ideal for ad-hoc queries. Automatic Clustering helps with data organization, but its impact on complex JSON queries might be limited. Using Query Acceleration alone requires a larger warehouse and may not address the underlying issue of unoptimized queries on semi-structured data. While a dedicated warehouse is a good practice, it does not address the underlying performance issue related to JSON queries.

NEW QUESTION # 61

You are implementing a data share between two Snowflake accounts. The provider account wants to grant the consumer account access to a function that returns anonymized customer data based on a complex algorithm. The provider wants to ensure that the consumer cannot see the underlying implementation details of the anonymization algorithm. Which of the following approaches can achieve this goal? (Select TWO)

- A. Create an external function in the provider account and grant usage to the share. Share the share with the consumer account.
- B. Share the underlying table and provide the consumer account with the anonymization algorithm separately.
- C. Create a standard UDF in the provider account and grant usage on the UDF to the share. Share the share with the consumer account.
- D. **Create a secure UDF in the provider account and grant usage on the secure UDF to the share. Share the share with the consumer account.**
- E. Create a view that calls the secure UDF and share that view with the consumer account.

Answer: D,E

Explanation:

A secure UDF hides the underlying implementation details from the consumer. Option 'A' achieves this directly. Creating a view (Option 'D') that calls the secure UDF provides another layer of abstraction, further protecting the algorithm's implementation. A standard UDF (Option B) does not hide the implementation. Sharing the table directly (Option C) defeats the purpose of anonymization. While external functions exist (Option E), they would be unnecessarily complex in this scenario, which can be achieved natively through secure UDF and View combination.

NEW QUESTION # 62

You are tasked with designing a data pipeline to load data from an Azure Blob Storage container into Snowflake using an external stage. The data is in CSV format, compressed using GZIP. The container contains millions of small CSV files. To optimize the data loading process and minimize cost, which of the following strategies would you implement, considering both stage configuration and COPY INTO options? Choose TWO that apply.

- A. **Leverage Snowflake's Snowpipe with a REST API endpoint to trigger data loads whenever new files are available in the Azure Blob Storage container.**
- B. Use the 'VALIDATION MODE = RETURN ERRORS' option in the 'COPY INTO' statement to identify and correct any

data quality issues during the load. This ensures that only clean data is loaded into Snowflake.

- C. Consolidate the small CSV files in the Azure Blob Storage container into larger files before loading them into Snowflake. This reduces the overhead of processing numerous small files.
- D. Use the 'MATCH BY COLUMN NAME = CASE INSENSITIVE' option with a copy transformation in the 'COPY INTO' statement to ensure that the column order in the CSV files doesn't affect the data load.
- E. Create a pipe object with 'AUTO INGEST = TRUE' to automatically ingest new files as they are added to the Azure Blob Storage container. This ensures near real-time data ingestion.

Answer: A,C

Explanation:

Consolidating small files into larger files (option E) significantly improves COPY INTO performance by reducing overhead. Using Snowpipe with a REST API endpoint (option C) allows for efficient, triggered data loading. 'VALIDATION MODE' (option A) is useful for data quality but doesn't directly address optimization for millions of small files. 'AUTO INGEST' (option B) is specific to AWS S3 and Google Cloud Storage, not Azure Blob Storage event notifications. Column matching (option D) addresses schema flexibility, not optimization.

NEW QUESTION # 63

You are tasked with implementing row-level filtering on a 'customers' table in Snowflake. You need to restrict access to customer data based on the user's region. The 'customers' table has a 'region' column, and you have a mapping table 'user_regions' that associates users with the regions they are allowed to access. The 'user_regions' table has columns 'username' and 'region'. Which of the following SQL statements correctly creates and applies a row access policy to achieve this, minimizing complexity and maximizing performance? Select all that apply:

- A. Option C
- B. Option A
- C. Option D
- D. Option B
- E. Option E

Answer: A,E

Explanation:

Options C and E are correct. They both create a row access policy that checks if the current role is ACCOUNTADMIN (C) or either ACCOUNTADMIN/SECURITYADMIN (E). If not, it checks if the current user is associated with the region in the 'user_regions' table. The policy is then applied to the 'customers' table on the 'region' column. Option A is incorrect because it only checks for 'admin_role' but ACCOUNTADMIN is needed. Option B only provides row access functionality and no additional administrative access. Option D is incorrect because it doesn't have administrative override. The OR REPLACE Clause is used to ensure there are no errors when running this for multiple times or running if exists scenarios.

NEW QUESTION # 64

You are tasked with optimizing a continuous data pipeline that loads data from an external stage into a Snowflake table using streams.

The pipeline is experiencing significant latency during peak hours. The stream is defined on a very large table with frequent updates and deletes. Which of the following strategies would be MOST effective in reducing the latency of the data pipeline, considering stream performance and cost implications?

- A. Implement a materialized view on top of the stream to pre-aggregate the data.
- B. Implement a more aggressive pruning strategy on the base table to reduce the amount of data that the stream needs to track.
- C. Create multiple streams on the same base table, each filtering for specific types of changes (e.g., INSERT, UPDATE, DELETE).
- D. Reduce the RETENTION TIME of the stream. This will limit the amount of historical data tracked and improve performance.
- E. Increase the size of the virtual warehouse used for loading data. This will provide more compute resources for processing the stream.

Answer: B

Explanation:

Pruning on the base table directly reduces the amount of data the stream has to track, therefore improving performance. Increasing warehouse size (A) might help but is not directly related to the stream's efficiency. Creating multiple streams (C) will likely increase the overhead. Reducing RETENTION _ TIME (D) might lead to data loss if changes are not consumed promptly. A materialized view (E) consumes credits and does not directly affect stream performance. The correct answer focuses on minimizing the workload for the stream itself.

NEW QUESTION # 65

We are well acknowledged for we have a fantastic advantage over other vendors - We offer you the simulation test with the Soft version of our DEA-C02 exam engine: in order to let you be familiar with the environment of DEA-C02 test as soon as possible. Under the help of the real simulation, you can have a good command of key points which are more likely to be tested in the real DEA-C02 test. Therefore that adds more confidence for you to make a full preparation of the upcoming DEA-C02 exam.

DEA-C02 Latest Exam Fee: <https://www.vceengine.com/DEA-C02-vce-test-engine.html>

They have utmost faith in our DEA-C02 products, After you pay we will send you download links, account and password of Snowflake DEA-C02 Latest Exam Fee exam dump materials in a minute, Snowflake Study DEA-C02 Group Less time for high efficiency with our exam materials, There are professional DEA-C02 latest dumps pdf and DEA-C02 exam dumps in VCEEngine, Snowflake Study DEA-C02 Group They often encounter situations in which the materials do not match the contents of the exam that make them waste a lot of time and effort.

After all, wasn't the world considered to be flat at one DEA-C02 point by the best minds of the time. You've probably thought about a number of ways you can do your job better.

They have utmost faith in our DEA-C02 products, After you pay we will send you download links, account and password of Snowflake exam dump materials in a minute.

Quiz Snowflake - Latest Study DEA-C02 Group

Less time for high efficiency with our exam materials, There are professional DEA-C02 latest dumps pdf and DEA-C02 exam dumps in VCEEngine, They often encounter situations in which the Latest DEA-C02 Test Report materials do not match the contents of the exam that make them waste a lot of time and effort.

What's more, part of that VCEEngine DEA-C02 dumps now are free: <https://drive.google.com/open?id=1o1aZQecCmpvgUbyLQEY56LSRiojXLFEG>