

DEA-C02 Braindumps Pdf & Valid Braindumps DEA-C02 Ebook



DOWNLOAD the newest Actual4test DEA-C02 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=16M9dxODN0dS5_9QgFsQZUHNmdKxEq6pu

Our website is a pioneer in providing comprehensive Snowflake dumps torrent because we have a group of dedicated IT experts who have more than 10 years of experience in the study of DEA-C02 test questions and answers. They work in advance to make sure that our candidates will get latest and accurate DEA-C02 Exam Prep materials. You will get DEA-C02 passing score with the shortest duration for exam preparation.

In order to allow our customers to better understand our DEA-C02 quiz prep, we will provide clues for customers to download in order to understand our DEA-C02 exam torrent in advance and see if our products are suitable for you. As long as you have questions, you can send us an email and we have staff responsible for ensuring 24-hour service to help you solve your problems. We do not charge extra service fees, but the service quality is high. Your satisfaction is the greatest affirmation for us and we sincerely serve you. Our DEA-C02 Exam Guide deliver the most important information in a simple, easy-to-understand language that you can learn efficiently learn with high quality. Whether you are a student or an in-service person, our DEA-C02 exam torrent can adapt to your needs.

>> **DEA-C02 Braindumps Pdf** <<

2026 Snowflake DEA-C02: Authoritative SnowPro Advanced: Data Engineer (DEA-C02) Braindumps Pdf

In Actual4test's website you can free download study guide, some exercises and answers about Snowflake Certification DEA-C02 Exam as an attempt.

Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Sample Questions (Q109-Q114):

NEW QUESTION # 109

A data engineer is using the Snowflake Spark connector to write data to a Snowflake table. The write operation fails consistently with the error 'net.snowflake.client.jdbc.SnowflakeSQLException: SQL execution error: String ' is too long (maximum is 16777216)'. Which of the following is the most likely cause and how can it be resolved using Spark Connector?

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

Answer: C

Explanation:

The error message clearly states that a string is too long. While options A might seem plausible, it is better to pre-process data on spark using 0, 16777216). This helps to ensure data conforms to the desired constraints before reaching snowflake. Option B addresses performance, but is not a root cause of the error. Option C relates to nulls but the error message is for maximum string length.

NEW QUESTION # 110

You are developing a data pipeline in Snowflake that uses SQL UDFs for data transformation. You need to define a UDF that calculates the Haversine distance between two geographical points (latitude and longitude). Performance is critical. Which of the following approaches would result in the most efficient UDF implementation, considering Snowflake's execution model?

- A. Create a SQL UDF leveraging Snowflake's VECTORIZED keyword, hoping to automatically leverage SIMD instructions, without any code changes to mathematical calculation inside the UDF
- B. Create a SQL UDF that directly calculates the Haversine distance using Snowflake's built-in mathematical functions (SIN, COS, ACOS, RADIANS). This is straightforward and easy to implement.
- C. Create a SQL UDF that pre-calculates the RADIANS for latitude and longitude only once and stores them in a temporary table, using those values for subsequent distance calculations within the same session.
- D. Create an External Function (using AWS Lambda or Azure Functions) to calculate the Haversine distance. This allows for offloading the computation to a separate compute environment.
- E. Create a Java UDF that calculates the Haversine distance, leveraging optimized mathematical libraries. This allows for potentially faster execution due to lower-level optimizations.

Answer: B

Explanation:

SQL UDFs are generally the most efficient for simple calculations within Snowflake because they are executed within the Snowflake engine, minimizing data movement and overhead. While Java UDFs (option B) can offer optimizations, the overhead of invoking the Java environment often outweighs the benefits for this type of calculation. External Functions (option C) introduce significant latency due to network communication. Option D provides temporary performance improvements for the specific session, but is not the most efficient general solution. Vectorized keyword doesn't exist in snowflake to create UDFs, Hence it won't allow compilation. This question's emphasis is on understanding the trade-offs between different UDF types and their performance implications within the Snowflake architecture.

NEW QUESTION # 111

Consider a scenario where you have a Snowflake table named 'CUSTOMER DATA' containing customer IDs (INTEGER) and encrypted credit card numbers (VARCHAR). You need to create a secure JavaScript UDF to decrypt these credit card numbers using a custom encryption key stored securely within Snowflake's internal stage, and then mask all but the last four digits of the decrypted number for data protection. Which of the following actions are necessary to ensure both functionality and security while adhering to Snowflake's best practices for UDF development and security?

- A. Store the encryption key in a separate file on an internal stage accessible only by the UDF's service account and load the key from the file within the UDF at runtime.
- B. Encrypt the key using a weaker encryption algorithm before storing it in an internal stage to balance security and performance.
- C. Store the encryption key directly within the JavaScript UDF code as a string variable.
- D. Pass the encryption key as an argument to the UDF each time it is called.
- E. Use Snowflake's Secure Vault (Secret) feature to store the encryption key and retrieve it securely within the UDF.

Answer: A,E

Explanation:

Options B and D are the correct answers. Option B - Storing the encryption key in a file on an internal stage, accessible only by the UDF's service account, is a secure way to manage the key. Option D - Snowflake's Secure Vault (Secret) feature is designed specifically for securely storing and managing sensitive information like encryption keys. This is the most recommended approach. Options A and C are insecure and should be avoided. Option E defeats the purpose of encryption.

NEW QUESTION # 112

Consider a table 'EVENT DATA' that stores events from various applications. The table has columns like 'EVENT ID', 'EVENT TIMESTAMP', 'APPLICATION ID', 'USER ID', and 'EVENT _ TYPE'. A significant portion of queries filter on 'EVENT TIMESTAMP' ranges AND 'APPLICATION ID'. The data volume is substantial, and query performance is crucial. You observe high clustering depth after initial loading. Which combination of actions will provide the MOST effective performance optimization, addressing both clustering depth and query performance?

- A. Cluster the table on 'EVENT TIMESTAMP' and periodically run 'OPTIMIZE TABLE EVENT DATA' using a small warehouse. Also, create a separate table clustered on 'APPLICATION'
- B. Cluster the table on '(EVENT TIMESTAMP, APPLICATION ID)' and periodically run 'OPTIMIZE TABLE EVENT DATA' using a warehouse sized appropriately for the table size. Then, monitor clustering depth regularly.
- C. Create separate tables for each ' ', each clustered on 'EVENT_TIMESTAMP'. Then, create a view that UNION ALLs these tables.
- D. Cluster the table on 'USER ID' and rely solely on Snowflake's automatic reclustering feature, without running 'OPTIMIZE TABLES' manually.
- E. Create multiple materialized views: one filtering on common 'EVENT TIMESTAMP' ranges, and another filtering on common 'APPLICATION ID' values.

Answer: B

Explanation:

Clustering on '(EVENT _ TIMESTAMP, APPLICATION_ID)' directly addresses the common query patterns. Regularly running 'OPTIMIZE TABLE EVENT DATA' with an appropriately sized warehouse ensures the data remains well-clustered as new data is added, reducing clustering depth and maintaining performance. Monitoring clustering depth is essential to identify when reclustering is needed. Clustering on a single dimension like 'USER IDS (C)' doesn't address the primary query patterns. Creating separate tables (A, D) introduces complexity and management overhead. Materialized views (E) are helpful for specific pre-aggregated results, but clustering optimizes the base table for a wider range of queries. Optimizing with the right sized warehouse is crucial, a small warehouse might take an extremely long time.

NEW QUESTION # 113

You have an external table in Snowflake pointing to data in Azure Blob Storage. The data consists of customer transactions, and new files are added to the Blob Storage daily. You want to ensure that Snowflake automatically picks up these new files and reflects them in the external table without manual intervention. However, you are observing delays in Snowflake detecting the new files. What are the potential reasons for this delay and how can you troubleshoot them? (Choose two)

- A. The Azure Event Grid notification integration is not properly configured to notify Snowflake about new file arrivals in the Blob Storage.
- B. Snowflake's internal cache is not properly configured; increasing the cache size will solve the problem.
- C. The file format used for the external table is incompatible with the data files in Blob Storage.
- D. The external table's 'AUTO_REFRESH' parameter is set to 'FALSE', which disables automatic metadata refresh.
- E. The storage integration associated with the external table does not have sufficient permissions to access the Blob Storage.

Answer: A,D

Explanation:

The two primary reasons for delays in Snowflake detecting new files in an external table are: 1) Incorrect configuration of the cloud provider's notification service (Azure Event Grid in this case). Snowflake relies on these notifications to be informed about new file arrivals. If the integration isn't set up correctly, Snowflake won't know when to refresh the metadata. 2) The parameter must be set to ' TRUE' for automatic metadata refresh to occur. If it's set to FALSE, manual refreshes are required using 'ALTER EXTERNAL TABLE ... REFRESH'. Options D and E, although possible issues, won't directly cause a delay in detecting new files, but rather cause issues accessing files after detection. Option C is irrelevant as Snowflake's caching mechanism does not directly impact external table metadata refresh.

NEW QUESTION # 114

.....

In order to help customers solve problems, our company always insist on putting them first and providing valued service. We deeply believe that our DEA-C02 question torrent will help you pass the exam and get your certification successfully in a short time. Maybe you cannot wait to understand our DEA-C02 Guide questions; we can promise that our products have a higher quality when

compared with other study materials. At the moment you can free download the demo of our DEA-C02 guide torrents, and I can make a bet that you will be fond of our DEA-C02 exam questions if you understand it.

Valid Braindumps DEA-C02 Ebook: https://www.actual4test.com/DEA-C02_examcollection.html

One of the most important functions of our DEA-C02 preparation questions are that can support almost all electronic equipment, including the computer, mobile phone and so on, Snowflake DEA-C02 Braindumps Pdf Plenty of concepts get mixed up together due to which student feel difficult to identify them, Snowflake DEA-C02 Braindumps Pdf You find us, you find the way to success.

Render design elements from Photoshop, Illustrator, DEA-C02 and other visual development environments into usable elements of a working website. With Photos, music, movies, email, and hundreds of Valid Braindumps DEA-C02 Ebook of thousands of apps, the iPad helps you organize, explore, and maximize your digital world.

Quiz 2026 Snowflake Valid DEA-C02 Braindumps Pdf

One of the most important functions of our DEA-C02 Preparation questions are that can support almost all electronic equipment, including the computer, mobile phone and so on.

Plenty of concepts get mixed up together due to which student feel Valid Braindumps DEA-C02 Ebook difficult to identify them, You find us, you find the way to success, We will provide you with the best quality exam materials.

So once you made the resolution to choose us, we will not let you down.

P.S. Free & New DEA-C02 dumps are available on Google Drive shared by Actual4test: <https://drive.google.com/open?id=16M9dxODNOdS59QgFsQZUHNmdKxEq6pu>