

高效率地使用我們可靠的DEA-C02最新考證: SnowPro Advanced: Data Engineer (DEA-C02), Snowflake DEA-C02考試很容易通過



順便提一下，可以從雲存儲中下載TestpdfDEA-C02考試題庫的完整版：<https://drive.google.com/open?id=1JHmTCSOOdNuQcMhfs80ZQNIDIDLnzMT>

如果你想在短時間內，以最小的努力，達到最有效果的結果Snowflake的DEA-C02考試準備，你可以使用Testpdf Snowflake的DEA-C02考試培訓資料，Testpdf的培訓資料是通過實踐檢驗了的，也是通過眾多考生證明了它確實可以百分百通過考試，利用了它，你將達到你的目的，得到最佳的效果。

在Testpdf的網站上你可以免費下載Testpdf為你提供的關於Snowflake DEA-C02 認證考試學習指南和部分練習題及答案作為嘗試。

>> DEA-C02最新考證 <<

DEA-C02考試備考經驗和知識點 - 通過DEA-C02認證考試最好的方法

作為一位 DEA-C02 考生而言，作好充分的準備可以幫助您通過考試。Testpdf的 DEA-C02 題庫覆蓋了最新的 DEA-C02 考試指南及考試真題題型。DEA-C02 隸屬於 Snowflake 認證考試科目。我們的 DEA-C02 認證考題已經幫助很多考生通過考試，試題質量和考題的覆蓋率都有保證，保證考生權利不受任何損失。獲取 DEA-C02 考試認證證書可以用來實施一些複雜多變的工程。

最新的 SnowPro Advanced DEA-C02 免費考試真題 (Q318-Q323):

問題 #318

You are responsible for monitoring data quality in a Snowflake data warehouse. Your team has identified a critical table, 'CUSTOMER DATA', where the 'EMAIL' column is frequently missing or contains invalid entries. You need to implement a solution that automatically detects and flags these anomalies. Which of the following approaches, or combination of approaches, would be MOST effective in proactively monitoring the data quality of the 'EMAIL' column?

- A. Create a Snowflake Task that executes a SQL query to count NULL 'EMAIL' values and invalid 'EMAIL' formats (using regular expressions). The task logs the results to a separate monitoring table and alerts the team if the count exceeds a predefined threshold.
- B. Use Snowflake's Data Quality features (if available) to define data quality rules for the 'EMAILS' column, specifying acceptable formats and thresholds for missing values. Configure alerts to be triggered when these rules are violated.
- C. Utilize an external data quality tool (e.g., Great Expectations, Deequ) to define and run data quality checks on the 'CUSTOMER DATA' table, integrating the results back into Snowflake for reporting and alerting.
- D. Implement a Streamlit application connected to Snowflake that visualizes the percentage of NULL and invalid 'EMAIL' values over time, allowing the team to manually monitor trends.
- E. Schedule a daily full refresh of the 'CUSTOMER DATA' table from the source system, overwriting any potentially corrupted data.

答案: A,B,C

解題說明：

Options A, B, and D are the most effective. Option A provides a programmatic approach within Snowflake. Option B leverages Snowflake's built-in data quality capabilities (if available, check documentation for supported features and editions). Option D integrates with external specialized tools. Option C relies on manual monitoring, which is less proactive. Option E does not address the root cause of data quality issues and could potentially overwrite valid data with erroneous data.

問題 #319

You are loading JSON data into a Snowflake table with a 'VARIANT' column. The JSON data contains nested arrays with varying depths. You need to extract specific values from the nested arrays and load them into separate columns in your Snowflake table. Which approach would provide the BEST performance and flexibility?

- A. Use Snowpipe with auto-ingest, loading directly into the table with the 'VARIANT' column. Define data quality checks with pre-load data transformation.
- B. Create a view with nested 'FLATTEN' functions to extract the values from the 'VARIANT' column. The view serves as the source for further transformations.
- C. Load the entire JSON into a 'VARIANT' column and then use SQL with nested 'FLATTEN' functions to extract the desired values during query time.
- D. Use a stored procedure to parse the JSON data and insert values into the table row by row.
- E. Use a 'COPY' command with a 'TRANSFORM' clause that uses JavaScript UDFs to parse the JSON and extract the values during the load process. Load the extracted values directly into the target columns.

答案： E

解題說明：

Using a 'COPY' command with a 'TRANSFORM' clause and JavaScript UDFs allows for efficient parsing and extraction of values during the load process. This minimizes the amount of data stored in the 'VARIANT' column and avoids expensive query-time parsing. Stored procedures perform row by row operations which are inefficient. Using Flatten functions could be useful to denormalise json, but javascript parsing during load is better. Snowpipe and auto-ingest just move the challenge to a real-time streaming scenario, which may not be optimized for transforming data into a relational structure.

問題 #320

A financial institution needs to mask sensitive customer data (PII) in a 'CUSTOMER' table. The table contains columns like 'CUSTOMER ID', 'FIRST NAME', 'LAST NAME', 'CREDIT CARD', and 'ADDRESS'. The data should be masked differently for different roles: 'ANALYST' role should see obfuscated values for names and addresses, while the 'SUPPORT' role should see the last four digits of the credit card and a hashed version of the address. The 'CUSTOMER ID' should never be masked. Assume a central masking policy already exists called 'PII MASKING POLICY'. Which of the following statements is the MOST efficient and secure way to achieve this?

- A. Create view for each role which applies masking functions to the columns. Grant SELECT access on those views to relevant roles.
- B. Create multiple masking policies, one for each role and sensitive column combination, each with the appropriate masking expression. Then, apply each masking policy individually to its respective column. Use the function to implement role-based masking within each policy.
- C. Create external functions to handle the complex masking logic and call them from the masking policy.
- D. Create a single masking policy with a complex stored procedure that checks the current role and applies different masking functions accordingly, then apply this policy to all sensitive columns.
- E. Create multiple masking policies with different masking expressions and apply them directly to the columns based on the role using conditional expressions within the policies. Use 'CASE' statements within the masking policy to differentiate between roles.

答案： E

解題說明：

Option A is the most efficient and recommended approach. Using conditional expressions (CASE statements or similar logic with within a single masking policy allows you to manage different masking rules for different roles within a centralized policy. This avoids creating numerous policies and simplifies management. Option B is less efficient due to the overhead of calling a stored procedure for each row access. Option C using view is difficult to maintain and not scalable. Option D can lead to a proliferation of masking policies, making management cumbersome. Option E introduces external dependencies and potential latency.

問題 #321

A data engineering team is tasked with optimizing a complex query that joins three tables: 'ORDERS', 'CUSTOMERS', and 'PRODUCTS'. The 'ORDERS' table contains millions of records and is frequently joined with 'CUSTOMERS' (containing customer demographics) and 'PRODUCTS' (containing product details). The initial query uses standard JOIN syntax, but performance is slow. The query retrieves order details along with customer and product information, filtering by a specific date range in the 'ORDERS' table and a customer segment in the 'CUSTOMERS' table. Which optimization strategy would be MOST effective for significantly improving query performance?

- A. Increase the virtual warehouse size to X-LARGE without analyzing the query profile.
- B. Apply clustering keys to the 'ORDERS' table based on the date column used in the WHERE clause and clustering keys to the 'CUSTOMERS' table on the customer segment column. Also create appropriate indexes.
- C. Convert the entire dataset into a single VARIANT column and query using JSON path expressions.
- D. Create materialized views that pre-join the 'ORDERS', 'CUSTOMERS', and 'PRODUCTS' tables and filter based on common criteria.
- E. Replace the standard JOINS with LATERAL FLATTEN operations.

答案： B,D

解題說明：

Materialized views can pre-compute the joined data and apply filters, significantly reducing the amount of data the query needs to process at runtime. Clustering keys on the frequently filtered columns in 'ORDERS' and 'CUSTOMERS' will enable data skipping, allowing Snowflake to efficiently prune micro-partitions. Creating indexes is not directly supported in Snowflake, however Clustering is equivalent to Indexing strategy. Increasing warehouse size without analysis is not an appropriate strategy. Lateral Flatten is relevant when dealing with array or semi- structured data, and converting the data into VARIANT would severely impact query performance.

問題 #322

A data engineer is tasked with creating a Snowpark Python UDF to perform sentiment analysis on customer reviews. The UDF, named 'analyze_sentiment', takes a string as input and returns a string indicating the sentiment ('Positive', 'Negative', or 'Neutral'). The engineer wants to leverage a pre-trained machine learning model stored in a Snowflake stage called 'models'. Which of the following code snippets correctly registers and uses this UDF?

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

答案： C

解題說明：

The most concise and recommended way to define a Snowpark UDF in Python is using the @F.udf decorator. This decorator automatically handles registration with Snowflake and simplifies the code. It also correctly specifies the 'return_type', 'input_types', and required packages'. Options A, B, C and E are either missing the decorator or have issues with specifying input types or session usage. The session.add_packageS is not a proper way to define packages used by UDFs and 'StringType' is not imported from 'snowflake.snowpark.types', so the correct way is to set return_type and input_types within the decorator.

問題 #323

.....

如果你是找考試資料或學習書籍？試試我們的免費的 Snowflake 的 DEA-C02 考題吧！這是一個免費試用考試PDF 測試版本的考題，你可以類比真實的考試情景，可以快速讓你掌握 Snowflake 的基礎知識。我們的 DEA-C02 權威考試題庫軟體是 Snowflake 認證廠商的授權產品。正確率100%，讓你一次性輕鬆通過 Snowflake DEA-C02 考試。

DEA-C02考試心得：<https://www.testpdf.net/DEA-C02.html>

如果你想在短時間內，以最小的努力，達到最有效果的結果，就來使用我們的 Snowflake SnowPro Advanced: Data

蕭峰有點無語，轉移話題道，他眼眸壹冷，也沒再去追，如果你想在短時間內，以最小的努力，達到最有效果的結果，就來使用我們的 Snowflake SnowPro Advanced: Data Engineer (DEA-C02) - DEA-C02 考古題培訓資料吧，SnowPro Advanced: Data Engineer (DEA-C02)考古題由資深的IT專家團隊研究出來的結果。

以下是最新的DEA-C02題庫資訊，由Testpd題庫網負責收集，DEA-C02題庫的高效率和準確性兩大特點讓我們收到廣大考生的好評，獲得如此有價值的認證方案對您來說是非常划算的，面對職場的競爭和不景氣時期，提升您的專業能力是未DEA-C02來最好的投資，獲得與自己所學、有興趣或工作相關SnowPro Advanced 證照、無論是升學、升職還是未來的職場競爭，比別人更具優勢！

- 此外，這些TestpdfDEA-C02考試題庫的部分內容現在是免費的：<https://drive.google.com/open?id=1JHmITCSOodNuOcMhfs80ZONlDlDlnzMT>

此外，這些TestpdfDEA-C02考試題庫的部分內容現在是免費的：<https://drive.google.com/open?id=1JHmITCSOodNuOcMhfs80ZONlDlDlnzMT>