

SOL-C01測試 - SOL-C01學習筆記



2026 PDFExamDumps最新的SOL-C01 PDF版考試題庫和SOL-C01考試問題和答案免費分享：https://drive.google.com/open?id=1DqDMgr_pJcPSo-G4GVtx9j1wLY3U78WJ

Snowflake SOL-C01 認證試題庫學習資料根據最新的知識點以及輔導資料進行整編，覆蓋面廣，蘊含了眾多最新的Snowflake 考試知識點。如果你正在準備 SOL-C01 考試並且像我一樣急需通過，那 SOL-C01 認證試題剛好可以幫助你。因為完善的 SOL-C01 學習資料資料覆蓋 Snowflake 考試所有知識點，減少你考試的時間成本和經濟成本，助你輕鬆通過考試

Snowflake SOL-C01 考試大綱：

主題	簡介
主題 1	• Data Loading and Virtual Warehouses: This domain covers loading structured, semi-structured, and unstructured data using stages and various methods, virtual warehouse configurations and scaling strategies, and Snowflake Cortex LLM functions for AI-powered operations.
主題 2	• Identity and Data Access Management: This domain focuses on Role-Based Access Control (RBAC) including role hierarchies and privileges, along with basic database administration tasks like creating objects, transferring ownership, and executing fundamental SQL commands.
主題 3	• Interacting with Snowflake and the Architecture: This domain covers Snowflake's elastic architecture, key user interfaces like Snowsight and Notebooks, and the object hierarchy including databases, schemas, tables, and views with practical navigation and code execution skills.
主題 4	• Data Protection and Data Sharing: This domain addresses continuous data protection through Time Travel and cloning, plus data collaboration capabilities via Snowflake Marketplace and private Data Exchange sharing.

>> SOL-C01測試 <<

SOL-C01學習筆記 - SOL-C01題庫

IT專業技術認證是進入IT行業的“敲門磚”。由國際著名IT企業頒發的職業證書，證明了你具有某種專業IT技能，為國際承認並通用。這些國際著名 IT企業為：Microsoft、Oracle、Cisco、Amazon、IBM、Oracle等。SOL-C01 考試就是其中一個流行的 Snowflake 認證。許多考生對這門考試沒有什麼信心，其實，SOL-C01 最新的擬真試題是用最快和最聰明的的方式來傳遞您的考試，並幫助您獲得 Snowflake SOL-C01 認證。

最新的 SnowPro Advanced SOL-C01 免費考試真題 (Q51-Q56):

問題 #51

You have a Snowflake external stage configured to access Parquet files in an AWS S3 bucket.

You want to query these Parquet files directly using Snowflake without explicitly loading them into a table. You also want to enable directory tables for this external stage. Which of the following steps are necessary to enable directory tables for the external stage and query the data?

- A. Create an external stage pointing to the S3 bucket with 'ENABLE DIRECTORY = TRUE' specified during stage creation, then no refresh is needed. Directory tables are automatically updated.
- B. Create an external stage pointing to the S3 bucket, then execute 'ALTER STAGE REFRESH;', then execute 'ALTER STAGE SET ENABLE_DIRECTORY = TRUE;'
- C. Create an external stage pointing to the S3 bucket, execute 'ALTER STAGE SET ENABLE_DIRECTORY = TRUE;', then refresh the directory table using 'ALTER STAGE REFRESH;'
- D. Create an external stage pointing to the S3 bucket, then execute 'ALTER STAGE SET ENABLE_DIRECTORY = TRUE;'
- E. Create an external stage pointing to the S3 bucket with 'ENABLE DIRECTORY = TRUE' specified during stage creation, then refresh the directory table using 'ALTER STAGE REFRESH;'

答案: E

解題說明:

To enable directory tables, you must set 'ENABLE DIRECTORY = TRUE' when creating the stage or using ALTER STAGE. Then, 'ALTER STAGE REFRESH;' populates the directory table with the file metadata. The directory table is not automatically updated and needs to be refreshed periodically.

問題 #52

How does Snowflake process queries?

- A. Using MPP compute clusters
- B. With shared-disk architecture
- C. Through third-party connectors
- D. By optimizing data in cloud storage

答案: A

解題說明:

Snowflake processes queries using Massively Parallel Processing (MPP) compute clusters, deployed as virtual warehouses. Each warehouse consists of multiple compute nodes working in parallel to execute queries efficiently. When a query is submitted, Snowflake distributes tasks across nodes, processes data subsets concurrently, and aggregates results. This architecture enables high performance, scalability, and the ability to handle complex analytical workloads. While Snowflake does incorporate elements of shared-disk storage, query execution itself depends on MPP compute clusters. Options such as third-party connectors or storage optimization do not represent the core query processing mechanism.

問題 #53

You are tasked with using the 'PARSE DOCUMENT' function in Snowflake to extract key information (name, address, phone number) from a large collection of scanned invoices stored as PDF files in an AWS S3 bucket. The invoices have varying formats and quality. Which of the following approaches would be MOST effective to structure the extracted data for analysis?

- A. Employ a combination of 'PARSE DOCUMENT' and Snowflake's external functions to integrate with a third-party OCR service for enhanced text extraction. Post-process the results using SQL to load structured data into a table. Use 'MAX FILE SIZE=16777216' while calling parse document
- B. Directly load PDF files into a relational table's TEXT column and write SQL queries utilizing LIKE and SUBSTRING functions to extract invoice details. Use 'MAX FILE SIZE=16777216' while loading pdf
- C. Use 'PARSE DOCUMENT' with a pre-defined JSON schema to enforce a rigid structure on the extracted data, and load the output directly into a relational table with predefined columns for name, address, and phone number.
- D. Create a custom UDF (User-Defined Function) that calls 'PARSE_DOCUMENT' and then uses regular expressions or other text processing techniques to extract the desired information from the raw text output. Load the extracted data into a relational table.
- E. Use 'PARSE DOCUMENT' with default settings and load the raw JSON output into a VARIANT column in a Snowflake table. Implement custom SQL queries to extract the desired information from the VARIANT data.

答案: D

解題說明:

Option C provides the most robust and flexible approach. Given the varying formats and quality of the invoices, a pre-defined JSON schema (option B) is unlikely to work effectively. Loading raw JSON into a VARIANT column (option A) requires extensive post-processing. Option D, while potentially effective, introduces the complexity and cost of a third-party OCR service. And MAX_FILE_SIZE parameter controls the maximum size, in bytes, of a single uncompressed file that can be loaded from the stage. Option E is not a scalable and efficient approach.

問題 #54

What is the typical syntax for using the LIMIT clause to retrieve the first 10 rows of a table named employees?

- A. SELECT * FROM employees WHERE LIMIT 10;
- B. SELECT * FROM employees LIMIT 10;
- C. SELECT LIMIT 10 * FROM employees;
- D. LIMIT 10 SELECT * FROM employees;

答案: B

解題說明:

The LIMIT clause is used to restrict the number of rows returned by a query. The correct syntax in Snowflake is: SELECT * FROM employees LIMIT 10;

This retrieves the first 10 rows of the result set. LIMIT appears at the end of the SELECT statement.

Snowflake follows a SQL syntax similar to MySQL/PostgreSQL regarding LIMIT usage.

Incorrect forms include:

- * Using LIMIT before SELECT
- * Embedding LIMIT inside WHERE
- * Reversing the order of SELECT and LIMIT

LIMIT is frequently used for sampling, record previewing, debugging, and exploratory analysis.

問題 #55

A Snowflake account uses separate virtual warehouses for different workloads: 'REPORTING_WH' for reporting queries and

'TRANSFORM_WH' for data transformation jobs. A user, Alice, is responsible for running both types of workloads. Alice is experiencing performance issues with reporting queries when the transformation jobs are running concurrently. Both warehouses are currently configured with the

'STANDARD' size. What are the recommended strategies to improve the performance of reporting queries without significantly impacting the performance of transformation jobs?

- A. Enable multi-cluster warehouses for 'REPORTING_WH' with a scaling policy set to 'STANDARD' and a minimum of 2 clusters. Implement Query Acceleration Service to reduce the compute resources used by reporting queries and implement Resource Monitors to limit the credit usage of 'TRANSFORM_WH'.
- B. Increase the size of 'REPORTING_WH' to 'LARGE' and enable auto-suspend on to prevent it from consuming resources when idle.
- C. Enable Resource Monitors to limit the credit usage of 'TRANSFORM_WH' and increase the size of to 'LARGE'.
- D. Enable multi-cluster warehouses for both and set the scaling policy to 'ECONOMY' for 'TRANSFORM_WH' and 'STANDARD' for 'REPORTING_WH'.
- E. Enable auto-suspend on 'REPORTING_WH' and increase the size of 'TRANSFORM_WH' to 'LARGE'.

答案: A

解題說明:

Enabling a multi-cluster warehouse for 'REPORTING_WH' with a 'STANDARD' scaling policy allows Snowflake to automatically scale up the number of clusters to handle the increased workload from concurrent reporting queries. The Query Acceleration Service (QAS) will try to reduce the resource consumption of reporting queries. Implementing Resource Monitors for 'TRANSFORM_WH' will prevent it from consuming excessive credits and impacting the performance of 'REPORTING_WH' during concurrent execution.

• • • • •

SOL-C01學習筆記: https://www.pdfexamdumps.com/SOL-C01_valid-braindumps.html

- P.S. PDFExamDumps在Google Drive上分享了免費的、最新的SQL-C01考試題庫：https://drive.google.com/open?id=1DqDMgr_pJcPSo-G4GVtx9j1wLY3U78WJ