

最新のSOL-C01参考書 & 資格試験のリーダープロバイダー & 更新のSOL-C01過去問題



2026年ShikenPASSの最新SOL-C01 PDFダンプおよびSOL-C01試験エンジンの無料共有: <https://drive.google.com/open?id=1NxOuUAoNHSw7D3vJzg540Lze4O6AILZ9>

ShikenPASSが提供した教育資料は真実のテストに非常に近くて、あなたが弊社の短期の特殊訓練問題を通じてすぐにSnowflake専門の知識を身につけられます。弊社は君のSOL-C01試験の100%合格率を保証いたします。

SOL-C01試験資格証明書を取得することは難しいです。でも、Snowflake SOL-C01復習教材を選べれば、試験に合格することは簡単です。SOL-C01復習教材の内容は全面的で、価格は合理的です。そして、Snowflakeはお客様にディスカウントコードを提供でき、SOL-C01復習教材をより安く購入できます。

>> SOL-C01参考書 <<

SOL-C01過去問題、SOL-C01認定試験

IT業種は急激に発展しているこの時代で、IT専門家を称賛しなければならないです。彼らは自身が持っている先端技術で色々な便利を作ってくれます。それに、会社に大量な人的・物的資源を節約させると同時に、案外の良い効果を取得しました。彼らの給料は言うまでもなく高いです。そのような人になりたいのですか。羨ましいですか。心配することはないです。ShikenPASSのSnowflakeのSOL-C01トレーニング資料はあなたに期待するものを与えますから。ShikenPASSを選ぶのは、成功を選ぶということになります。

Snowflake SOL-C01 認定試験の出題範囲:

トピック	出題範囲
トピック 1	• 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.
トピック 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.
トピック 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.
トピック 4	• 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.

Snowflake Certified SnowPro Associate - Platform Certification 認定 SOL-C01 試験問題 (Q196-Q201):

質問 # 196

A data engineer needs to grant SELECT privileges on all tables within a newly created schema named 'SALES DATA' to a role called 'ANALYST ROLE'. However, new tables may be added to this schema in the future. What is the most efficient and secure way to achieve this, ensuring that future tables automatically inherit the necessary permissions?

- A. Grant the USAGE privilege on the database containing the schema and SELECT privilege on 'FUTURE GRANTS for tables in 'SALES DATA' to the 'ANALYST ROLE using 'GRANT USAGE ON DATABASE TO ROLE ANALYST ROLE; GRANT SELECT ON FUTURE TABLES IN SCHEMA SALES DATA TO ROLE ANALYST ROLE;'
- B. Grant SELECT privilege to 'ANALYST ROLE directly on each table using 'GRANT SELECT ON TABLE ... TO ROLE ANALYST ROLE;' after table creation.
- C. Create a custom role that inherits from 'ANALYST ROLE and grant 'SELECT privilege on all tables to this new role.
- D. Use a stored procedure to automatically grant the SELECT privilege to 'ANALYST ROLE whenever a new table is created in the 'SALES DATA' schema.
- E. Grant 'SELECT privilege to the PUBLIC role on all tables within the 'SALES_DATX schema.

正解: A

解説:

Using GRANT SELECT ON FUTURE TABLES' ensures that any new tables created in the schema automatically inherit the SELECT privilege for the specified role. Option A requires manual intervention for each new table. Options C and E are unnecessarily complex. Option D poses a security risk by granting privileges to the PUBLIC role.

質問 # 197

You have enabled directory tables on an external stage pointing to an AWS S3 bucket. The S3 bucket contains millions of small JSON files. When you query the directory table, you notice that the query performance is slow, even when filtering by 'METADATA\$FILENAME'. What strategies can you implement to optimize the performance of querying the directory table?

- A. Create an index on the 'METADATA\$FILENAME column of the directory table.
- B. Periodically refresh the directory table using 'ALTER STAGE REFRESH;'.- C. Increase the virtual warehouse size used for querying the directory table.
- D. Create a materialized view on top of the directory table, including the 'METADATA\$FILENAME , , and any other frequently filtered columns.
- E. Reduce the number of files in the S3 bucket. Directory tables perform better with fewer files.

正解: D

解説:

Creating a materialized view (Option B) is the most effective way to optimize the performance of querying a directory table. Materialized views pre-compute and store the results of a query, allowing for faster retrieval of data, especially when the underlying table (the directory table in this case) is large. Increasing warehouse size (Option A) can help to a certain extent, but a materialized view provides a more significant performance boost. Refreshing the directory table (Option C) keeps the metadata up-to-date but doesn't directly improve query performance.

Creating an index (Option D) is not possible on directory tables. While reducing the number of files (Option E) might help slightly, it's often impractical and doesn't address the underlying issue of querying a large dataset. It's best to handle millions of files with proper indexing and data partitioning.

質問 # 198

A junior data analyst accidentally deleted a critical table using Snowsight. After restoring the table from Time Travel, they want to prevent this from happening again. What is the MOST appropriate preventative measure they should take using Snowsight, considering they still need access to query other tables?

- A. Configure multi-factor authentication for the analyst's Snowflake account.
- B. Revoke the DELETE privilege on all tables for the analyst.
- C. Revoke the DELETE privilege ONLY on the specific critical table.

- D. Grant the analyst the MONITOR privilege on the database.
- E. Revoke the ACCOUNTADMIN role from the analyst.

正解: C

解説:

Revoking the DELETE privilege ONLY on the specific critical table prevents accidental deletion of that table while still allowing the analyst to delete data from other tables where they have the appropriate permissions. Revoking DELETE on all tables or removing ACCOUNTADMIN is too restrictive, as the analyst needs to query and potentially modify other tables. Multi-factor authentication enhances security but doesn't directly prevent accidental deletions. MONITOR privilege is for observing, not preventing actions.

質問 # 199

In the Query Profile, what does the Pruning section provide?

- A. Information on how Snowflake removed rows from the query results.
- B. Information on how Snowflake removed columns from the query results.
- C. Information on how Snowflake removed micro-partitions from the query scan.
- D. Information on how Snowflake removed objects from the query plan.

正解: C

解説:

The Pruning section of the Snowsight Query Profile shows how Snowflake eliminated unnecessary micro-partitions from the scan phase of the query. Snowflake stores data in micro-partitions and maintains metadata such as min/max values for each column within each partition. When a query includes filters (e.g., WHERE clauses), Snowflake evaluates this metadata to determine which micro-partitions cannot possibly satisfy the predicate. These partitions are skipped, meaning they are never scanned or read from storage. This process drastically improves performance because Snowflake minimizes I/O, reduces compute usage, and shortens execution time. Partition pruning is especially impactful on large tables because only a fraction of the stored micro-partitions typically need to be accessed.

The Pruning section does not show removed rows—that happens during the filter step. It does not show removed columns—column pruning is handled separately by the optimizer. It also does not show removed objects from the plan. Its sole purpose is to document micro-partition elimination and scan reduction.

質問 # 200

What is the primary use of Snowflake Cortex LLM functions?

- A. Optimizing data loading performance.
- B. Analyzing text data for sentiment summarization and other AI/ML tasks.
- C. Running machine learning models on structured data.
- D. Performing traditional SQL joins.

正解: B

解説:

Cortex LLM functions integrate large language model capabilities directly inside Snowflake. These functions enable:

- * Sentiment analysis
- * Summarization
- * Document question answering
- * Text classification
- * Embedding generation for vector search

They operate directly on Snowflake data, require no external infrastructure, and support secure in-platform AI workloads.

質問 # 201

.....

私たちに知られているように、適切な学習計画はすべての人々にとって非常に重要です。競争力を高めるために、学習計画を立てる必要があります。SOL-C01の実際の試験は、優れた学習計画の作成に役立つと考えています。SOL-C01学習教材を使用して、限られた時間でモデルテストを行うことができます。モデルテストを完

さらに、ShikenPASS SOL-C01ダンプの一部が現在無料で提供されています: <https://drive.google.com/open?id=1NxOuUAoNHSw7D3vJzg540Lz4O6AILZ9>