

SOL-C01 Pass Guarantee, Latest SOL-C01 Demo



2026 Latest GetValidTest SOL-C01 PDF Dumps and SOL-C01 Exam Engine Free Share: <https://drive.google.com/open?id=1cL6Jcol7AYY0OQTsdvZHAmRfW9dmbNWH>

Our valid Snowflake SOL-C01 dumps make the preparation easier for you. With these real SOL-C01 Questions, you can prepare for the test while sitting on a couch in your lounge. Whether you are at home or traveling anywhere, you can do SOL-C01 exam preparation with our Snowflake SOL-C01 Dumps. Snowflake Certified SnowPro Associate - Platform Certification (SOL-C01) test candidates with different learning needs can use our three formats to meet their needs and prepare for SOL-C01 test successfully in one go. Read on to check out the features of these three formats.

Snowflake SOL-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">• 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.
Topic 2	<ul style="list-style-type: none">• 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.
Topic 3	<ul style="list-style-type: none">• 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.
Topic 4	<ul style="list-style-type: none">• 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.

Latest SOL-C01 Demo & Test SOL-C01 Guide Online

While making revisions and modifications to the Snowflake SOL-C01 practice exam, our team takes reports from over 90,000 professionals worldwide to make the Snowflake Certified SnowPro Associate - Platform Certification exam questions foolproof. To make you capable of preparing for the Snowflake SOL-C01 Exam smoothly, we provide actual Snowflake SOL-C01 exam dumps.

Snowflake Certified SnowPro Associate - Platform Certification Sample Questions (Q51-Q56):

NEW QUESTION # 51

A Snowflake table 'product catalog' exists with columns 'product id', 'product_name', and 'price'.

You have a pipe-delimited file in an external stage named 's3 stage' containing product data, but some rows have missing values for the 'price' column. You want to load the data, replacing missing prices with a default value of 0. Which approach is the MOST appropriate to handle missing values during the 'COPY INTO' operation?

- A. The 'COPY INTO' command cannot handle missing values; pre-processing is always required.
- B. Load the data into a staging table, then use a SQL 'UPDATE' statement to replace null prices with 0.
- **C. Use a transformation within the 'COPY INTO' statement to replace null values with 0.**
- D. Modify the file format options to automatically replace null values with 0.
- E. Use the 'ON_ERROR = SKIP_FILE' option to skip any file with missing prices.

Answer: C

Explanation:

Using a transformation within the 'COPY INTO' statement (option B) allows you to directly handle missing values during the data loading process, providing the most efficient solution. You can utilize a 'CASE' statement or similar logic within the column mapping of the 'COPY INTO' statement to replace nulls with the default value of 0. Option A is not valid. Option C requires additional steps. Option D skips over valuable data, and Option E is false because of the transformation support.

NEW QUESTION # 52

You have a base table called 'customer_data' which contains sensitive information like credit card numbers. You need to create a view called that exposes only non-sensitive columns (e.g., customer_id, customer_name, city) and hides the logic of the underlying table's structure.

Additionally, you want to ensure that no user can directly query the 'customer_data' table. Which of the following steps are necessary to achieve this?

- **A. Create a secure view 'customer_summary' selecting only the desired columns from 'customer_data' .
Revoke SELECT privilege on 'customer_data' from all roles except the role that owns the table.**
- B. Create a materialized view 'customer_summary' selecting only the desired columns from 'customer_data'. Revoke SELECT privilege on 'customer_data' from all roles except the role that owns the table.
- C. Create a standard view 'customer_summary' selecting only the desired columns from 'customer_data' . Grant SELECT privilege on to the required roles. Do not revoke privileges on the base table.
- D. Create a standard view 'customer_summary' selecting only the desired columns from 'customer_data'. Revoke SELECT privilege on 'customer_data' from all roles except the role that owns the table.
- E. Create a secure view 'customer_summary' selecting all columns from 'customer_data'. Grant SELECT privilege on 'customer_summary' to the required roles.

Answer: A

Explanation:

Option C provides the correct solution. Creating a secure view 'customer_summary' hides the base table structure from users and ensures that they can only access the data through the view.

Revoking SELECT privileges on the 'customer_data' table prevents direct access to the sensitive data, making it accessible only through the view, and only to those with privilege on the view.

Materialized views are not appropriate here because their data is physically stored, and the need is to hide the base table not improving the performance. Option D is incorrect because without revoking the base table privilege, users can query the base table directly. Option E will display all the columns.

NEW QUESTION # 53

You are working with a very large table 'TRANSACTIONS' and need to improve the performance of queries that filter data based on a specific range of transaction timestamps. Which of the following is the MOST appropriate Snowflake feature to optimize these queries, and why?

- A. Repartitioning the table daily to ensure even data distribution across micro-partitions.
- B. Creating a standard B-tree index on the transaction timestamp column.
- C. Using a Snowflake Sequence to generate sequential transaction IDs, as this will automatically improve query performance.
- **D. Creating a Search Optimization Service on the table, as this can significantly accelerate point lookup and range queries.**
- E. Converting the table to a transient table to reduce metadata overhead.

Answer: D

Explanation:

The Search Optimization Service (SOS) is specifically designed to accelerate point lookup and range queries on large tables. SOS automatically creates and manages search access paths based on common query patterns, making it significantly more efficient than a standard index, which Snowflake doesn't directly support. Sequences don't inherently improve query performance. Snowflake doesn't support manual repartitioning. Transient tables affect data recovery, not query performance.

NEW QUESTION # 54

What information can be obtained by describing a table in Snowflake? (Choose any 3 options)

- **A. Table constraints**
- B. Table Origin
- **C. Indexes and keys**
- **D. Column names and data types**

Answer: A,C,D

Explanation:

The DESCRIBE TABLE (or DESC TABLE) command in Snowflake provides core metadata about table structure. This includes column names, data types, nullability, and default values. It also returns constraints if applied, such as PRIMARY KEY, UNIQUE, and CHECK constraints. Snowflake additionally displays information regarding clustering keys and other table characteristics. The command does not show table origin; Snowflake does not track lineage information directly through DESCRIBE TABLE. Indexes in Snowflake are not traditional B-tree indexes-Snowflake uses micro-partition pruning instead-but DESCRIBE TABLE can show clustering keys, which function similarly by enabling optimized data skipping. Therefore, column definitions, constraints, and keys are all valid outputs of DESCRIBE TABLE.

NEW QUESTION # 55

What is a primary function of a view in Snowflake?

- A. Managing user authentication
- **B. Providing a virtual table based on the result set of a query**
- C. Storing raw data
- D. Executing network configurations

Answer: B

Explanation:

A view in Snowflake represents a virtual table whose contents are defined by a stored SQL query. Views allow users to encapsulate transformation logic, simplify complex joins, enforce column-level security, and provide curated datasets for downstream consumers. Views do not store data themselves; the underlying tables store the data, and Snowflake executes the view's query each time the view is referenced. This makes views ideal for abstraction layers, semantic modeling, and separating compute costs across user groups. Incorrect options:

