

Latest Snowflake SOL-C01 Test Blueprint - Dump SOL-C01 File



Snowflake SOL-C01 SnowPro Associate - Platform Certification

**Questions & Answers PDF
(Demo Version – Limited Content)**

For More Information – Visit link below:

<https://p2pexam.com/>

Visit us at: <https://p2pexam.com/sol-c01>

What's more, part of that Actual4dump SOL-C01 dumps now are free: <https://drive.google.com/open?id=1arBO3RCqaChh-P5sDpiw5yseHoWa-ZWk>

Since the software keeps a record of your attempts, you can overcome mistakes before the SOL-C01 final exam attempt. Knowing the style of the Snowflake SOL-C01 examination is a great help to pass the test and this feature is one of the perks you will get in the desktop practice exam software.

The Actual4dump is one of the top-rated and renowned platforms that has been offering real and valid Snowflake Certified SnowPro Associate - Platform Certification (SOL-C01) exam practice test questions for many years. During this long time period countless Snowflake Certified SnowPro Associate - Platform Certification (SOL-C01) exam candidates have passed their dream SOL-C01 certification and they are now certified Snowflake professionals and pursuing a rewarding career in the market.

>> **Latest Snowflake SOL-C01 Test Blueprint** <<

Latest Snowflake Certified SnowPro Associate - Platform Certification pass review & SOL-C01 getfreedumps study materials

We provide the Snowflake SOL-C01 exam questions in a variety of formats, including a web-based practice test, desktop practice exam software, and downloadable PDF files. Actual4dump provides proprietary preparation guides for the certification exam offered by the SOL-C01 Exam Dumps. In addition to containing numerous questions similar to the SOL-C01 exam, the Snowflake Certified SnowPro Associate - Platform Certification (SOL-C01) exam questions are a great way to prepare for the Snowflake SOL-C01 exam dumps.

Snowflake SOL-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	<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 2	<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 3	<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 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.

Snowflake Certified SnowPro Associate - Platform Certification Sample Questions (Q109-Q114):

NEW QUESTION # 109

When querying semi-structured data in Snowflake, which function extracts a specific field from a JSON document?

- A. `PARSE_JSON()`
- B. `SELECT_FIELD()`
- C. `GET_PATH()`
- D. `EXTRACT_JSON()`

Answer: A

NEW QUESTION # 110

You are using Snowsight to analyze query performance. You notice a query that frequently accesses a table 'orders' is slow. Which of the following Snowsight features and actions could you use to identify potential bottlenecks and improve query performance? Select all that apply.

- A. Use Snowsight's data lineage feature to trace the origin of the 'orders' table and identify any upstream transformations that might be causing performance issues.
- B. Examine the query profile in Snowsight to identify stages with high execution time and resource consumption.
- C. Review the SQL text formatting in the query history to ensure it follows Snowflake's best practices.
- D. Use the 'Advisor' feature in Snowsight to receive recommendations on potential performance improvements, such as creating indexes or optimizing table clustering.
- E. Disable query history in Snowsight to reduce the overhead on the system.

Answer: B,D

Explanation:

The query profile in Snowsight provides detailed information about each stage of a query's execution, allowing you to identify bottlenecks. The 'Advisor' feature offers recommendations for performance improvements. Data lineage (Option B) might be helpful in some cases, but it's not the most direct way to address query performance. Disabling query history (Option C) is counterproductive for performance analysis. Option E has no direct impact on query performance. Only relevant functions in Snowsight related to Snowflake are acceptable.

NEW QUESTION # 111

What is the Snowsight Query Profile used for?

- A. To create new database objects
- **B. To visualize and analyze query performance**
- C. To execute SQL queries
- D. To manage data loading processes

Answer: B

Explanation:

The Snowsight Query Profile is a powerful diagnostic tool that provides a visual breakdown of how Snowflake executed a query. Its primary purpose is to help users visualize and analyze query performance. It displays execution steps, including scan operations, join strategies, pruning results, aggregation methods, and data movement between processing nodes.

The profile shows metrics such as execution time per step, partition pruning effectiveness, bytes scanned, and operator relationships. This allows developers, analysts, and DBAs to identify bottlenecks—such as unnecessary full-table scans, non-selective filters, or inefficient joins—and tune SQL accordingly.

Query Profile does not execute queries; execution happens in worksheets or programmatic interfaces. It does not create objects or manage data loading; those tasks involve separate SQL commands and UI interfaces.

Overall, Query Profile is essential for performance tuning, helping teams reduce compute costs, optimize warehouse sizing, and improve query efficiency.

NEW QUESTION # 112

What is the purpose of auto-resume on a virtual warehouse?

- A. To automatically scale up the warehouse size.
- **B. To automatically start the warehouse when a query is submitted to it.**
- C. To restart the warehouse manually after suspension.
- D. To prevent the warehouse from suspending.

Answer: B

Explanation:

When auto-resume is enabled, Snowflake automatically starts a suspended warehouse when a query or task is submitted. This ensures smooth operations without manual intervention.

It does not scale the warehouse, prevent suspension, or require manual restarting.

NEW QUESTION # 113

A data engineer is tasked with loading JSON files containing customer reviews from an external stage into a Snowflake table. The JSON files have varying schemas and nested structures.

Which of the following methods is the MOST efficient and scalable way to ingest and query this data, minimizing the need for upfront schema definition?

- A. Load the JSON data into a cloud storage location, then use a Spark cluster to process and transform the data into a structured format before loading it into Snowflake.
- B. Create a relational table with predefined columns based on the expected schema. Use a COPY INTO statement to load the data, handling schema variations through error handling in the COPY INTO statement.
- C. Use a stored procedure to parse the JSON files and insert the data into a relational table with a predefined schema. Implement error handling within the stored procedure to handle schema variations.
- **D. Create a VARIANT column in a Snowflake table. Load the JSON data into the VARIANT column using a COPY INTO statement. Query the data using dot notation and FLATTEN function as needed.**
- E. Create an external table using the CREATE EXTERNAL TABLE command, pointing to the external stage. Define a schema on the external table based on a representative JSON file. Query the data directly from the external table.

Answer: D

Explanation:

Using a VARIANT column is the most efficient way to load JSON data with varying schemas because it allows you to load the data as is without defining a rigid schema upfront. Dot notation and the FLATTEN function allow you to query the data flexibly. Creating a relational table (Option A) requires defining a schema upfront which is not ideal for varying schemas. External tables (Option C)

