

# SOL-C01 Exam Bootcamp & SOL-C01 Latest Dumps & SOL-C01 Study Materials



BONUS!!! Download part of PrepAwayExam SOL-C01 dumps for free: [https://drive.google.com/open?id=1cYUs7m5hS\\_hh-wQ3jhwKY6ZTDXvY7KHI](https://drive.google.com/open?id=1cYUs7m5hS_hh-wQ3jhwKY6ZTDXvY7KHI)

Snowflake SOL-C01 gives practice material that is as per the legitimate Snowflake SOL-C01 exam. A free demo is other than open to test the parts prior to buying the entire thing for the Snowflake SOL-C01. You can pass Snowflake Certified SnowPro Associate - Platform Certification on the off chance that you use Snowflake SOL-C01 Dumps material. Not withstanding zeroing in on our material, expecting that you went after in the Snowflake SOL-C01 exam, you can guarantee your cash back as per systems.

## Snowflake SOL-C01 Exam Syllabus Topics:

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

>> SOL-C01 Latest Exam Papers <<

## New SOL-C01 Exam Testking - SOL-C01 Valid Test Pdf

We are constantly updating our practice material to ensure that you receive the latest preparation material based on the actual Snowflake SOL-C01 exam content. Up to 1 year of free Snowflake Certified SnowPro Associate - Platform Certification (SOL-C01) exam questions updates are also available at PrepAwayExam. The PrepAwayExam offers a money-back guarantee (terms and conditions apply) for students who fail to pass their Snowflake Certified SnowPro Associate - Platform Certification (SOL-C01) exam on the first try.

## Snowflake Certified SnowPro Associate - Platform Certification Sample

## Questions (Q49-Q54):

### NEW QUESTION # 49

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 refresh the directory table using 'ALTER STAGE REFRESH;'.
- B. 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.
- 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, then execute 'ALTER STAGE REFRESH;', then execute 'ALTER STAGE SET ENABLE\_DIRECTORY = TRUE;'.

**Answer: A**

Explanation:

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.

### NEW QUESTION # 50

You are loading data from a CSV file stored in an Amazon S3 bucket into a Snowflake table named 'CUSTOMER DATA'. The CSV file contains a header row, and the data is comma-separated. The 'CUSTOMER DATA' table has columns 'customer\_id', 'first\_name', and 'email'.

You want to use a named file format object called 'CSV FORMAT' which you have already created. You also want to skip the header row and only load data where the column is not null.

Which of the following 'COPY INTO' statement achieves this most efficiently and correctly?

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

**Answer: E**

Explanation:

Option D is the most correct because it efficiently copies the data from the stage, using a named file format, skips the header and sets error handling. The WHERE clause to filter null emails cannot be directly used within COPY INTO. Options A, B and C, is incorrect, as 'WHERE clause and 'SKIP\_HEADER = 1' is not applicable as the way written in the option. E is incorrect as its using select statement and it cannot be used in 'COPY INTO'

### NEW QUESTION # 51

How does Snowflake's compute layer handle query execution?

- A. Using MPP (massively parallel processing) compute
- B. With shared-disk architecture
- C. By optimizing data in cloud storage
- D. Using single-threaded processing

**Answer: A**

Explanation:

Snowflake's compute layer uses Massively Parallel Processing (MPP), meaning queries are divided into smaller tasks distributed across multiple compute nodes in the Virtual Warehouse. Each node processes a portion of the data simultaneously, maximizing parallelism and drastically reducing query times.

Although Snowflake uses a central storage layer (shared-disk model), the compute engine behaves like a shared-nothing MPP system, where each node handles local processing independently, minimizing contention.

Incorrect options:

- \* Snowflake does not rely on single-thread execution.
- \* Storage optimization occurs at the Storage Layer, not compute.
- \* Snowflake does not use traditional shared-disk execution; compute nodes work in parallel independently.

This architecture enables high performance for large analytical workloads.

### NEW QUESTION # 52

A data engineer is tasked with creating a new database called 'SALES DATA' within a Snowflake account. They want to ensure that only users with the 'ACCOUNTADMIN' role or a custom role

'SALES ADMIN' can manage the database's overall configuration. Which of the following steps are necessary to achieve this? (Choose all that apply)

- A. Grant the 'MODIFY' privilege on the 'DATABASE' object to the 'SALES\_ADMIN' role.
- **B. Grant the 'OWNERSHIP' privilege on the 'SALES\_DATA' database to the 'SALES\_ADMIN' role.**
- C. Grant the 'CREATE SCHEMA' privilege on the 'SALES DATA' database to the 'SALES\_ADMIN' role.
- D. Grant the 'USAGE' privilege on the 'DATABASE' object to the 'SALES ADMIN' role.
- **E. The 'ACCOUNTADMIN' role already has sufficient privileges; no additional grants are needed for it.**

**Answer: B,E**

Explanation:

The 'OWNERSHIP' privilege grants the 'SALES ADMIN' role full control over the 'SALES DATA' database, including configuration. 'ACCOUNTADMIN' inherently has ownership over all objects and requires no additional grants. 'CREATE SCHEMA' is for creating schemas within the database, 'MODIFY' is not a valid privilege on 'DATABASE' object, and 'USAGE' allows access but not management of the database object itself.

### NEW QUESTION # 53

You have a Snowflake table named 'SALES DATA' that you load data into daily from a CSV file using Snowsight. Recently, the load times have increased significantly. You suspect the Virtual Warehouse size is the bottleneck. You have the following Virtual Warehouse sizes available: X-Small, Small, Medium, Large, X-Large. Choosing the best size depends on cost and speed. You examine the Query History in Snowsight and notice that the COPY INTO commands are consistently using only a small fraction of the X-Large warehouse's compute resources. Which of the following actions would be the MOST cost-effective while also potentially improving (or at least maintaining) the data loading performance?

- A. Increase the Virtual Warehouse size to 2X-Large. This will guarantee faster load times.
- B. Enable auto-suspend on the X-Large warehouse if it isn't already enabled. This will reduce costs when the warehouse is idle but won't address the loading bottleneck.
- C. Decrease the Virtual Warehouse size to Small or Medium. Since the current warehouse is underutilized, a smaller size will reduce costs without significantly impacting performance.
- D. Switch to using Snowpipe for continuous data ingestion instead of COPY INTO. This may improve load times but requires additional configuration and cost analysis.
- **E. Keep the X-Large warehouse but investigate other potential bottlenecks, such as the file format definition, network latency, or concurrency issues with other queries. Consider using a dedicated warehouse for loading only. Also consider increasing number of parallel threads for loading.**

**Answer: E**

Explanation:

Since the X-Large warehouse is underutilized, increasing its size is not cost-effective. Reducing the size might further degrade performance. Enabling auto-suspend only addresses idle costs.

The best approach is to investigate other potential bottlenecks while maintaining the current (potentially oversized) warehouse. Using a dedicated warehouse isolates the loading workload.

Snowpipe is a good alternative, but it's a more complex solution that requires evaluation and isn't the immediately cost-effective solution. Option D addresses both performance and cost concerns without making assumptions about Snowpipe's suitability in this

