

# Enhance Your Success Rate with DumpsActual's Snowflake DEA-C02 Exam Dumps



What's more, part of that DumpsActual DEA-C02 dumps now are free: <https://drive.google.com/open?id=1NXr7-QKiV4Bt8WLvReLvDDrYMzaxTUAU>

With our wide range of Snowflake DEA-C02 exam questions types and difficulty levels, you can tailor your Snowflake DEA-C02 exam practice to your needs. Your performance and exam skills will be improved with our Snowflake DEA-C02 Practice Test software. The software provides you with a range of Snowflake DEA-C02 exam dumps, all of which are based on past Snowflake DEA-C02 certifications.

Our Snowflake DEA-C02 practice exam simulator mirrors the DEA-C02 exam experience, so you know what to anticipate on DEA-C02 certification exam day. Our SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) practice test software features various question styles and levels, so you can customize your Snowflake DEA-C02 exam questions preparation to meet your needs.

>> [Pass4sure DEA-C02 Study Materials](#) <<

## New DEA-C02 Test Vce | Valid DEA-C02 Exam Prep

Our DEA-C02 study quiz boosts many advantages and it is your best choice to prepare for the test. Our DEA-C02 learning prep is compiled by our first-rate expert team and linked closely with the real exam. And our DEA-C02 training materials provide three versions and multiple functions to make the learners have no learning obstacles. The passing rate of our DEA-C02 Guide materials is high and you don't need to worry that you have spent money but can't pass the test.

## Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Sample Questions (Q262-Q267):

### NEW QUESTION # 262

You have created an external table in Snowflake that points to a large dataset stored in Azure Blob Storage. The data consists of JSON files, and you've noticed that query performance is slow. Analyzing the query profile, you see that Snowflake is scanning a large number of unnecessary files. Which of the following strategies could you implement to significantly improve query performance against this external table?

- A. Create a materialized view on top of the external table to pre-aggregate the data.
- **B. Convert the JSON files to Parquet format and recreate the external table to point to the Parquet files.**
- C. Increase the size of the Snowflake virtual warehouse to provide more processing power.
- D. Create an internal stage, copy all JSON Files, create and load the target table, and drop external table
- **E. Partition the data in Azure Blob Storage based on a relevant column (e.g., date) and define partitioning metadata in the external table definition using PARTITION BY.**

**Answer: B,E**

Explanation:

Partitioning the data (B) allows Snowflake to prune unnecessary files during query execution, significantly improving performance. Converting to Parquet (C) provides a columnar storage format that is more efficient for analytical queries compared to JSON, reducing I/O and processing time. Increasing warehouse size (A) might help but is not the most effective strategy. Materialized views (D) are not directly applicable to external tables. Copying all files and creating internal tables is not using the external table functionality (E).

### NEW QUESTION # 263

You are tasked with implementing row-level security (RLS) on a 'SALES' table to restrict access based on the 'REGION' column. Users with the 'NORTH REGION ROLE' should only see data where 'REGION = 'NORTH''. You've created a row access policy named 'north\_region\_policy'. After applying the policy to the 'SALES' table, users with the 'NORTH REGION ROLE' are still seeing all rows.

Which of the following is the MOST likely reason for this and how can it be corrected?

- A. The user has not logged out and back in since the role was granted to them. Force the user to re-authenticate.
- B. The policy function within is not using the correct context function to determine the user's role. It should use 'CURRENT\_ROLE()' instead of 'CURRENT\_USER()'
- C. The policy needs to be explicitly refreshed. Execute 'REFRESH ROW ACCESS POLICY north\_region\_policy ON SALES;'
- **D. The ' does not have the USAGE privilege on the database and schema containing the 'SALES' table. Grant the USAGE privilege to the role.**
- E. The is not enabled. Execute 'ALTER ROW ACCESS POLICY ON SALES SET ENABLED = TRUE;'

**Answer: D**

Explanation:

Row access policies require the role to have USAGE privilege on the database and schema. Without this privilege, the policy cannot be enforced. The other options, while potentially relevant in other scenarios, are not the most likely cause for the described issue. Row access policies are automatically enabled when applied and the correct context function would be CURRENT\_ROLE(). A refresh command is not required.

### NEW QUESTION # 264

A Snowflake data warehouse contains a table 'WEB EVENTS' with columns like 'EVENT ID', 'EVENT TIMESTAMP', 'USER', 'PAGE URL', and 'SESSION ID'. The data engineering team has enabled search optimization on 'PAGE URL' because analysts frequently filter on specific URLs. However, they notice that queries filtering on multiple 'PAGE URL' values (e.g., using 'WHERE PAGE URL IN ('url1', 'url2)', are not performing as well as expected. What are the potential reasons for this behavior, and what strategies can be used to improve performance in this scenario? Select all that apply:

- A. Search optimization is automatically disabled when using IN clause, therefore it is important to rewrite the query without using IN operator.
- B. Statistics on the 'PAGE URL' column are outdated. Run 'ANALYZE TABLE WEB EVENTS' to refresh the statistics.
- **C. Search optimization is not designed to efficiently handle IN list lookups with a large number of values. Consider using a temporary table or common table expression (CTE) to pre-filter the data.**
- D. The warehouse size is too small to handle the complexity of the IN list lookup. Increase the warehouse size.
- **E. The number of distinct values in the 'PAGE URL' column is very high, leading to a large search access path, making IN list lookups inefficient. Consider clustering by PAGE\_URL**

**Answer: C,E**

Explanation:

Options A and B are correct. A high cardinality in (Option A) means the search access path is large, making IN list lookups expensive. Search optimization is generally optimized for point lookups or small range scans, not large IN lists. Option B suggests using a temporary table or CTE to pre-filter the data, which can significantly improve performance by reducing the amount of data the optimizer needs to consider during the final filtering step. Option C (increasing warehouse size) might improve performance, but it doesn't directly address the issue of IN list lookups. Option D (analyzing the table) is always a good practice, but it's unlikely to be the primary cause of the slow IN list performance. Option E is wrong. Search optimization is not automatically disabled when using IN clause

#### NEW QUESTION # 265

You are tasked with designing a solution to load semi-structured data (JSON) from an AWS S3 bucket into a Snowflake table using Snowpipe and the REST API. The data in S3 is constantly being updated, and you need to ensure that only new or modified files are loaded into Snowflake. Which of the following steps are essential for implementing an efficient and cost-effective solution?

- A. Create a Snowflake external function that polls the S3 bucket every minute, checks for new files using the LIST command, and then calls the Snowpipe REST API endpoint for each new file.
- B. Use the 'VALIDATION MODES copy option with 'RETURN\_ALL RESULTS = TRUE to validate all data being loaded into the Snowflake table.
- C. Configure auto-ingest using SQS queue and SNOWPIPE object. No need to manually call the REST API endpoint for data loading.
- D. Configure Snowpipe to automatically detect new files in the S3 bucket using event notifications, but manually refresh the pipe using SYSTEM \$PIPE STATUS periodically to ensure that all files are processed.
- E. Configure an S3 event notification to trigger a REST API call to the Snowpipe endpoint whenever a new or modified file is added to the S3 bucket. The API call should include the file name in the request.

**Answer: C,E**

Explanation:

Options A and E are the most efficient and cost-effective solutions. Option A utilizes S3 event notifications to trigger Snowpipe, loading only new files using REST API, and avoids unnecessary polling. Option E uses Snowflake's Auto Ingest feature. Auto ingest eliminates the need for manual Snowpipe calls through REST API and reduces latency. Option B involves inefficient polling. Option C involves unnecessary manual refreshing of the pipe. Option D focuses on data validation during the copy process but doesn't address the core requirement of efficient file detection and triggering.

#### NEW QUESTION # 266

Consider the following Snowflake UDTF definition written in Python:

□ Which of the following statements are TRUE regarding the deployment and usage of this UDTF?

- A. The UDTF will automatically be available in all schemas across all databases in the Snowflake account.
- B. The UDTF needs to be registered using 'session.udtf.register' or 'create or replace function' with the 'imports' clause referencing the Python file, and the handler' specifying the function name.
- C. The UDTF can be called directly in SQL using 'SELECT FROM TABLE(process\_json(VARIANT COLUMN));' without any prior registration.
- D. The library needs to be explicitly installed and configured within the UDTF's environment using a Snowpark session.
- E. The return type of the generator 'yield' must strictly adhere to the declared output schema , or errors will occur during execution.

**Answer: B,E**

Explanation:

UDTFs in Snowflake require explicit registration using 'session.udtf.register' or 'create or replace function' command, defining the location of the source code (Python file) and the function to be executed. Also, data types of values which is 'yield'ed in the body of UDTF must strictly adhere with Output schema. Libraries from 'snowflake.snowpark' are usually available and does not needs explicit configuration. UDTFs are schema-bound, not automatically available everywhere. Direct call to UDTF without creation isn't possible.

#### NEW QUESTION # 267

.....



myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt,  
myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt,  
myportal.utt.edu.tt, bicyclebuysell.com, www.stes.tyc.edu.tw, learn.anantlibrary.in, myportal.utt.edu.tt, myportal.utt.edu.tt,  
myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt,  
myportal.utt.edu.tt, myportal.utt.edu.tt, Disposable vapes

DOWNLOAD the newest DumpsActual DEA-C02 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1NXr7-QKiV4Bt8WLvReLvDDrYMzaxTUAU>