

# Exam DEA-C02 Training, Dumps DEA-C02 Reviews



**DVA-C02**  
**AWS Certified**  
**Developer**  
**Associate**

**Certification Questions**  
**& Exams Dumps**

[www.edurely.com](http://www.edurely.com)

DOWNLOAD the newest ExamBoosts DEA-C02 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1UIH-yXLOX22TC6cw9qgB5bbIWbBIYgfs>

The users of our DEA-C02 exam questions log on to their account on the platform, at the same time to choose what they want to attend the exam simulation questions, the DEA-C02 exam questions are automatically for the user presents the same as the actual test environment simulation DEA-C02 test system, the software built-in timer function can help users better control over time, so as to achieve the systematic, keep up, as well as to improve the user's speed to solve the problem from the side with our DEA-C02 test guide.

We have created a number of reports and learning functions for evaluating your proficiency for the DEA-C02 exam dumps. In preparation, you can optimize Snowflake DEA-C02 practice exam time and question type by utilizing our Snowflake DEA-C02 Practice Test software. ExamBoosts makes it easy to download SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) exam questions immediately after purchase.

>> [Exam DEA-C02 Training](#) <<

## Valid DEA-C02 exam training material & cost-effective DEA-C02 PDF files

For years our team has built a top-ranking brand with mighty and main which bears a high reputation both at home and abroad. The sales volume of the DEA-C02 test practice guide we sell has far exceeded the same industry and favorable rate about our DEA-C02 learning guide is approximate to 100%. Why the clients speak highly of our DEA-C02 reliable exam torrent? Our dedicated service, high quality and passing rate and diversified functions contribute greatly to the high prestige of our DEA-C02 exam questions.

## Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Sample Questions (Q195-Q200):

### NEW QUESTION # 195

You are developing a Secure UDF in Snowflake to encrypt sensitive customer data'. The UDF should only be accessible by authorized roles. Which of the following steps are essential to properly secure the UDF?

- A. Granting the EXECUTE privilege on the UDF only to the roles that require access.
- B. Using the 'SECURE' keyword when creating the UDF to prevent viewing the UDF definition.
- C. Setting the 'SECURITY INVOKER' clause when creating the UDF to execute the UDF with the privileges of the caller.
- D. Using masking policies instead of Secure UDFs is the recommended approach for data security
- E. Ensuring that the UDF is owned by a role with appropriate permissions and limiting access to this role.

**Answer: A,B,E**

Explanation:

Secure UDFs protect the code definition. Granting EXECUTE privilege controls access to the UDF Ownership control is critical for managing permissions. SECURITY INVOKER, when used inappropriately can lead to security breaches if not properly managed and it executes with the privileges of the caller, potentially bypassing intended access restrictions. Masking policies are useful, but don't cover the core security functionality of secure UDFs, which hide the function's code itself.

#### NEW QUESTION # 196

A Snowflake data warehouse contains a table named 'SALES TRANSACTIONS' with the following columns: 'TRANSACTION ID', 'PRODUCT ID', 'CUSTOMER ID', 'TRANSACTION DATE', and 'SALES AMOUNT'. You need to optimize a query that calculates the total sales amount per product for a given month. The 'SALES TRANSACTIONS' table is very large (billions of rows), and queries are slow. Given the following initial query: `SELECT PRODUCT ID, SUM(SALES AMOUNT) AS TOTAL SALES FROM SALES TRANSACTIONS WHERE TRANSACTION DATE BETWEEN '2023-01-07' AND '2023-01-31' GROUP BY PRODUCT ID;` Which of the following actions, when combined, would MOST effectively improve the performance of this query?

- A. Create a temporary table with the results of the query and query that table instead.
- B. Convert the column to a VARCHAR data type.
- C. Create a materialized view that pre-aggregates the total sales amount per product and month.
- D. Increase the virtual warehouse size to the largest available size.
- E. Create a clustering key on 'PRODUCT\_ID' and 'TRANSACTION\_DATE' columns in the 'SALES\_TRANSACTIONS' table.

**Answer: C,E**

Explanation:

Creating a clustering key on 'PRODUCT ID' and 'TRANSACTION DATE' allows Snowflake to efficiently prune micro-partitions based on the date range filter, and then quickly group by 'PRODUCT\_ID'. A materialized view pre-aggregates the data, significantly reducing the amount of computation required at query time. While increasing the warehouse size might provide some improvement, it is not the most efficient solution. Converting 'TRANSACTION\_DATE' to VARCHAR is detrimental. Using a temporary table is not necessarily an optimization.

#### NEW QUESTION # 197

A Snowflake table 'ORDERS' is clustered on the 'ORDER DATE' column. After several months, you notice that many micro-partitions contain data from a wide range of 'ORDER DATE' values, and query performance on date range filters is degrading. Which of the following actions could improve performance and reduce the overlap in micro-partitions?

- A. Run 'OPTIMIZE TABLE ORDERS' to recluster the table.
- B. Change the ORDER DATE column to VARCHAR and then cluster.
- C. Change the clustering key to a composite key including ORDER DATE and another relevant column.
- D. Drop and recreate the table with a different clustering key.
- E. Increase the virtual warehouse size used for loading data into the 'ORDERS' table.

**Answer: A,C**

Explanation:

Running 'OPTIMIZE TABLE ORDERS' will recluster the table based on the existing clustering key, potentially improving the organization of micro-partitions and reducing overlap. Changing the clustering key to a composite key that includes 'ORDER\_DATE' along with another relevant column that exhibits different characteristics (e.g., a high-cardinality column related to the order) might also help to create more distinct micro-partitions. While the virtual warehouse size (D) can affect load performance, it doesn't directly address the micro-partition overlap issue. Dropping and recreating the table (B) is a drastic measure and might not be necessary. Changing the data type of the 'ORDER DATE' column to VARCHAR (C) is generally not recommended and will likely

hinder query performance.

### NEW QUESTION # 198

You have a large Snowflake table 'WEB EVENTS' that stores website event data'. This table is clustered on the 'EVENT TIMESTAMP' column. You've noticed that certain queries filtering on a specific 'USER ID' are slow, even though 'EVENT TIMESTAMP' clustering should be helping. You decide to investigate further. Which of the following actions would be MOST effective in diagnosing whether the clustering on 'EVENT TIMESTAMP' is actually benefiting these slow queries?

- A. Query the 'QUERY\_HISTORY' view to see the execution time of the slow query and compare it to the average execution time of similar queries without a 'USER' filter.
- B. Run 'SYSTEM\$ESTIMATE QUERY COST' to estimate the query cost to see if the clustering is impacting the cost.
- C. Execute 'SHOW TABLES' and check the 'clustering\_key' column to ensure that the table is indeed clustered on 'EVENT\_TIMESTAMP'.
- D. Use the 'SYSTEM\$CLUSTERING\_INFORMATION' function to get the 'average\_overlaps' for the table and 'EVENT\_TIMESTAMP' column. A low value indicates good clustering.
- E. Run 'EXPLAIN' on the slow query and examine the 'partitionsTotal' and 'partitionsScanned' values. A significant difference indicates effective clustering.

### Answer: E

Explanation:

The 'EXPLAIN' command provides detailed information about the query execution plan. By examining the 'partitionsTotal' and 'partitionsScanned' values, you can directly see how many micro-partitions Snowflake considered vs. how many it actually scanned. A large difference suggests that the clustering is effectively pruning partitions based on the 'EVENT\_TIMESTAMP' filter. While 'SYSTEM\$CLUSTERING\_INFORMATION' provides a general overview of clustering quality, it doesn't tell you how it's performing for a specific query. Looking at query history or checking that the clustering key is defined is useful for verifying basic setup but doesn't directly diagnose the effectiveness for slow queries.

### NEW QUESTION # 199

You are designing a data pipeline using Snowpipe to ingest data from multiple S3 buckets into a single Snowflake table. Each S3 bucket represents a different data source and contains files in JSON format. You want to use Snowpipe's auto-ingest feature and a single Snowpipe object for all buckets to simplify management and reduce overhead. However, each data source has a different JSON schema. How can you best achieve this goal while ensuring data is loaded correctly and efficiently into the target table?

- A. Use a single Snowpipe with a generic FILE FORMAT that can handle all possible JSON schemas. Implement a VIEW on top of the target table to transform and restructure the data based on the source bucket.
- B. Use a single Snowpipe and leverage Snowflake's VARIANT data type to store the raw JSON data. Create separate external tables, each pointing to a specific S3 bucket, and use SQL queries to transform and load the data into the target table.
- C. Create a separate Snowpipe for each S3 bucket. Although this creates more Snowpipe objects, it allows you to specify a different FILE FORMAT and transformation logic for each data source.
- D. Since Snowpipe cannot handle multiple schemas with a single pipe, pre-process the data in S3 using an AWS Lambda function to transform all files into a common schema before they are ingested by the Snowpipe.
- E. Use a single Snowpipe and leverage Snowflake's ability to call a user-defined function (UDF) within the 'COPY INTO' statement to transform the data based on the S3 bucket path. The UDF can parse the bucket path and apply the appropriate JSON schema transformation.

### Answer: E

Explanation:

The most efficient and manageable approach is to use a single Snowpipe with a UDF to handle schema variations. The UDF can inspect the S3 bucket path (available as metadata within the 'COPY INTO' statement) and apply the correct transformation logic for each data source. Creating separate Snowpipes (A) adds unnecessary overhead. Using a generic 'FILE FORMAT' and a VIEW (B) might work for simple transformations, but it becomes complex with significant schema differences. Using VARIANT and external tables (C) defeats the purpose of Snowpipe. Pre-processing in S3 (D) adds complexity outside of Snowflake. UDF provides schema flexibility during ingest and leverages Snowpipe's capabilities directly.

## NEW QUESTION # 200

What does it mean to win a competition? Users of our DEA-C02 actual exam can give you good answers. They have improved their strength and proved their strength. Now they have more opportunities and they have the right to choose. Of course, the effective learning methods they learned during the use of our DEA-C02 Preparation materials also greatly enhanced their work. All of them had praised that our DEA-C02 exam questions are the best choice they had made to buy. So what are you waiting for? Just rush to buy our DEA-C02 practice guide!

**Dumps DEA-C02 Reviews:** <https://www.examboosts.com/Snowflake/DEA-C02-practice-exam-dumps.html>

I believe that you will be more inclined to choose a good service product, such as DEA-C02 learning question, Haste is waste: Making hasty decisions can cost you your money and result in DEA-C02 Exam, Furthermore if we have the updated version, our system will send the latest DEA-C02 exam dumps to your email address automatically, you don't need to worry about missing the latest version, you just need to concentrate your attention on practicing, and we will do the rest for you, Snowflake Exam DEA-C02 Training Although some of the hard copy materials contain mock examination papers, they do not have the automatic timekeeping system.

Setting a Default Route, Agile development processes foster better collaboration, innovation, and results, I believe that you will be more inclined to choose a good service product, such as DEA-C02 learning question.

## ExamBoosts Snowflake DEA-C02 Different Formats

Haste is waste: Making hasty decisions can cost you your money and result in DEA-C02 Exam. Furthermore if we have the updated version, our system will send the latest DEA-C02 exam dumps to your email address automatically, you don't need to worry DEA-C02 about missing the latest version, you just need to concentrate your attention on practicing, and we will do the rest for you.

Although some of the hard copy materials contain mock Reliable DEA-C02 Exam Dumps examination papers, they do not have the automatic timekeeping system, Excellent after sale service.

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

yXLOX22TC6cw9qgB5bbIWbBIYgfs