

# Snowflake DEA-C02 Exam Questions - Proven Way Of Quick Preparation



P.S. Free 2025 Snowflake DEA-C02 dumps are available on Google Drive shared by NewPassLeader: [https://drive.google.com/open?id=1zcbkD08YCLfMgYSNuBuThC\\_IA08wrC35](https://drive.google.com/open?id=1zcbkD08YCLfMgYSNuBuThC_IA08wrC35)

People always do things that will benefit them, so as get a certificate of the DEA-C02 test dumps. Obtaining a certificate means more opportunity, a good job, a better salary, and a bright. The benefits are numerous, and we give you a quicker method to achieve this. Our DEA-C02 Questions and answers list the knowledge point for you, and you just need to speed some of your time to practice. We are pass guarantee and money back guarantee. And the pass rate is 98.

Each product has a trial version and our products are without exception, literally means that our DEA-C02 guide torrent can provide you with a free demo when you browse our website of DEA-C02 prep guide, and we believe it is a good way for our customers to have a better understanding about our products in advance. Moreover if you have a taste ahead of schedule, you can consider whether our DEA-C02 Exam Torrent is suitable to you or not, thus making the best choice. What's more, if you become our regular customers, you can enjoy more membership discount and preferential services.

>> DEA-C02 Valid Exam Pattern <<

## Latest DEA-C02 Demo & DEA-C02 Exam Pass4sure

Our company has dedicated ourselves to develop the DEA-C02 study materials for all candidates to pass the exam easier, also has made great achievement after more than ten years' development. As the certification has been of great value, a right DEA-C02 study material can be your strong forward momentum to help you pass the exam like a hot knife through butter. On the contrary, it might be time-consuming and tired to prepare for the DEA-C02 Exam without a specialist study material. So it's would be the best decision to choose our DEA-C02 study materials as your learning partner.

## Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Sample Questions (Q317-Q322):

### NEW QUESTION # 317

A data engineer is tasked with processing a large dataset of customer orders using Snowpark Python. The dataset contains a column stored as a string in 'YYYY-MM-DD HH:MI:SS' format. They need to create a new DataFrame with only the orders placed in the month of January 2023. Which of the following code snippets achieves this most efficiently, considering potential data volume and query performance?

- A. `df.filter(month(to_date(col('order_date'))) = 1 and year(to_date(col('order_date'))) = 2023)`
- B. `df.filter(to_date(col('order_date')).month == 1 and year(to_date(col('order_date'))) == 2023)`
- C. `df.filter(to_date(col('order_date')).month == 1).filter(year(to_date(col('order_date'))) == 2023)`
- D. `df.filter(col('order_date').startswith('2023-01'))`

- E. 

```
df.filter(col('order_date').like('2023-01%'))
```

**Answer: C**

Explanation:

Option C is the most efficient. While all the options might produce the correct result, Option C directly converts the string to a DATE type and uses the 'months' and 'year' functions. Option A, B and E are inefficient as 'like', 'substring' and 'startswith' will not leverage date-related optimizations. Option D has a similar approach as C, but creates two filtering steps instead of one with 'to date', which could impact performance.

### NEW QUESTION # 318

A data engineer accidentally truncated a critical table 'ORDERS' in the 'SALES DB' database. The table contained important historical order data, and the data retention period is set to the default. Which of the following options represents the MOST efficient and reliable way to recover the truncated table and its data, minimizing downtime and potential data loss?

- A. Restore the entire Snowflake account to a previous point in time before the table was truncated.
- B. Use the UNDROP TABLE command to restore the table. If UNDROP fails, clone the entire SALES\_DB database to a point in time before the truncation using Time Travel.
- C. Create a new table 'ORDERS' and manually re-insert the data from the application's logs and backups.
- **D. Use Time Travel to create a clone of the truncated table from a point in time before the truncation. Then, swap the original table with the cloned table.**
- E. Contact Snowflake support and request them to restore the table from a system-level backup.

**Answer: D**

Explanation:

Option D is the most efficient and reliable. Cloning the table using Time Travel to a point before the truncation allows quick recovery with minimal data loss. The clone can then replace the truncated table. Option A relies on Snowflake support, which can be slow. Option B, UNDROP TABLE command, if the data retention period has passed or data has been purged then we cannot use it. Option C is manual and error-prone. Option E is an extreme measure and impacts the entire account.

### NEW QUESTION # 319

You are managing a Snowflake environment where data retention is set to the default 1 day for all databases and tables. You need to clone a production table, 'CUSTOMER DATA', to a development environment to test some complex transformations. However, after cloning, you realize that the original 'CUSTOMER DATA' table in production was accidentally dropped 2 days ago. Which of the following statements accurately describe the situation and what can be done, if anything?

- A. The cloned 'CUSTOMER DATA' table will only remain accessible if the source database was also cloned at the same time. If only the table was cloned, the cloned table is unusable after 1 day.
- B. The cloned 'CUSTOMER DATA' table in the development environment will remain accessible because the clone operation created a point-in-time snapshot.
- C. The cloned 'CUSTOMER DATA' table in the development environment will inherit the data retention period from the original 'CUSTOMER DATA' table, so it will also be inaccessible after 1 day if dropped in development.
- D. The cloned 'CUSTOMER DATA' table in the development environment will no longer be accessible because the source table has been dropped and the data retention period has expired. The clone is effectively useless.
- **E. While the source 'CUSTOMER DATA' is dropped and outside the retention period, the clone created will remain available and fully functional for querying. It is an independent copy. The clone is usable indefinitely, within the limits of your Snowflake account storage.**

**Answer: E**

Explanation:

Cloning creates a separate, independent copy of the table's metadata and data at the time of the clone operation. Even if the original table is dropped and the data retention period has passed for the original table, the cloned table remains accessible and fully functional, as it has its own independent copy of the data. Option B is incorrect as future changes to DEV will be recoverable based on DEV retention settings which can be independent from PROD.

### NEW QUESTION # 320

A data engineer is tasked with creating an external table that points to a directory in AWS S3 containing CSV files. The files have a header row and are comma-delimited. The engineer executes the following DDL statement:

- A. The statement will succeed if a file format named 'csv\_format' with the specified properties already exists.
- B. The statement will fail because the 'LOCATION' is missing the protocol (e.g., 's3://P').
- **C. The statement will fail because the 'SKIP HEADER' property is invalid and the 'skip\_header' parameter should be used in the 'CREATE FILE FORMAT' statement.**
- D. The statement will fail because the FIELD\_DELIMITER property should be 'DELIMITER'.
- E. The statement will fail because external tables cannot be created directly from CSV files without defining a stage.

**Answer: C**

Explanation:

The 'SKIP\_HEADER' property is deprecated in the 'CREATE EXTERNAL TABLE' statement. You must define this property in the file format. A stage is not strictly necessary for creating an external table, the location can be specified directly if the proper access permissions are in place. The 'FIELD\_DELIMITER' property is valid. Also if a File Format is already created correctly, it does not mean this statement will succeed, the parameters passed on the statement itself are relevant. So statement will fail because the 'SKIP\_HEADER' property is invalid and the 'skip\_header' parameter should be used in the 'CREATE FILE FORMAT' statement.

### NEW QUESTION # 321

Consider a scenario where you have a large dataset of sensor readings stored in a Snowflake table called 'SENSOR DATA'. You need to build an external function to perform complex calculations on these readings using a custom Python library hosted on AWS Lambda. The calculation requires significant computational resources, and you want to optimize the data transfer between Snowflake and the Lambda function. The following SQL is provided: `CREATE OR REPLACE EXTERNAL FUNCTION ARRAY) RETURNS ARRAY VOLATILE MAX BATCH ROWS = 2000 RETURNS NULL ON NULL INPUT API INTEGRATION = aws_lambda_integration AS 'arn:aws:lambda:us-east-1:123456789012:function:sensorProcessor'`; Which of the following options would further optimize the performance and reduce data transfer costs, assuming the underlying Lambda function is correctly configured and functional?

- A. Convert the input data to a binary format (e.g., using 'TO\_BINARY' and 'FROM\_BINARY' functions in Snowflake) before sending it to the Lambda function, and decode it in Lambda to reduce the size of the data being transmitted.
- B. Rewrite the custom Python library in Java and create a Snowflake User-Defined Function (UDF) instead of using an external function.
- **C. Reduce the number of columns passed to the external function by performing pre-aggregation or filtering on the data within Snowflake before calling the function.**
- **D. Compress the data before sending it to the external function and decompress it within the Lambda function. Update the Lambda function to compress the array of results before sending it back to Snowflake and use Snowflake+s functions to decompress it.**
- **E. Increase the 'MAX BATCH ROWS' parameter to the maximum allowed value to send larger batches of data to the external function. Ensure Lambda function memory is increased appropriately.**

**Answer: C,D,E**

Explanation:

The correct answers are A, B, and C. Option A reduces the amount of data transferred over the network, improving performance and reducing costs. Option B minimizes data transfer by sending only necessary data. Option C improves throughput by processing more rows per Lambda invocation, potentially reducing overall execution time. Option D requires a binary format compatible with both Snowflake and Lambda, which can be complex to implement and may not always provide significant benefits. Option E could improve performance by executing directly within Snowflake, but requires re-writing the code and may not be feasible if the Python library relies on specific dependencies not available in the Snowflake Java UDF environment.

### NEW QUESTION # 322

.....

We hope that you have understood the major features of our three formats. Now let's discuss the benefits you can get upon buying our SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) exam material today. The first benefit you can get is the affordable price. Our SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) practice material is not expensive and every applicant can

