

# Snowflake DEA-C02 New Practice Questions, Exam Sample DEA-C02 Questions

---

## Snowflake DEA-C02 Exam

Snowflake SnowPro Advanced - Data Engineer

<https://www.passquestion.com/dea-c02.html>



Pass Snowflake DEA-C02 Exam with PassQuestion DEA-C02 questions and answers in the first attempt.

<https://www.passquestion.com/>

---

1/4

BONUS!!! Download part of ValidVCE DEA-C02 dumps for free: [https://drive.google.com/open?id=1a23JSSvHty\\_3E1bw63vtUcdCB7PUtlE](https://drive.google.com/open?id=1a23JSSvHty_3E1bw63vtUcdCB7PUtlE)

For candidates who are going to buy the DEA-C02 training materials online, the safety of the website is significant. We have professional technicians examine the website every day, if you buying DEA-C02 exam braindumps from us, we will provide you with a clean and safe online shopping environment. Besides, we offer you free update for one year, and you can get the latest information about DEA-C02 Exam Braindumps timely, so that you can change learning ways according to the new changes.

Our Company is always striving to develop not only our DEA-C02 latest practice materials, but also our service because we know they are the aces in the hole to prolong our career. Reliable service makes it easier to get oriented to the DEA-C02 exam. The combination of DEA-C02 Exam Guide and sweet service is a winning combination for our company, so you can totally believe that we are sincerely hope you can pass the DEA-C02 exam, and we will always provide you help and solutions with pleasure, please contact us through email then.

>> Snowflake DEA-C02 New Practice Questions <<

## Pass Guaranteed Snowflake - Fantastic DEA-C02 New Practice Questions

When you follow with our DEA-C02 exam questions to prepare for your coming exam, you will deeply touched by the high-quality

and high-efficiency. Carefully devised by the professionals who have an extensive research of the DEA-C02 exam and its requirements, our DEA-C02 study braindumps are a real feast for all the candidates. And if you want to have an experience with our DEA-C02 learning guide, you can free download the demos on our website.

## Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Sample Questions (Q71-Q76):

### NEW QUESTION # 71

You have configured a Snowpipe to load data from an AWS S3 bucket into a Snowflake table. The data in S3 is updated frequently. You've noticed that despite the Snowpipe being active and the S3 event notifications being configured correctly, some newly added files are not being picked up by the Snowpipe. You run 'SYSTEM\$PIPE and see the 'executionstate' is 'RUNNING' but the 'pendingFileCount' remains at 0, even after new files are placed in the S3 bucket. Choose all of the reasons that could explain the observations.

- A. The SQS queue or SNS topic associated with the S3 event notifications has a message retention period that is too short. Messages containing event details for new files are being deleted before Snowpipe can process them.
- B. The file format specified in the Snowpipe definition does not match the actual format of the files being placed in the S3 bucket.
- C. The S3 event notification configuration is missing the 's3:ObjectCreated:' event type, meaning that new file creation events are not being sent to the SQS queue or SNS topic.
- D. There is an insufficient warehouse size configured for the Snowpipe. Increase the warehouse size for optimal performance.
- E. The IAM role associated with your Snowflake account does not have sufficient permissions to read from the S3 bucket. Specifically, it lacks the 's3:GetObject' permission.

**Answer: B,C,E**

Explanation:

The problem states that the snowpipe is active, but 'pendingFileCount' remains zero. This means that the events are not making its way to snowpipe. Insufficient permissions (A) on the IAM role will prevent Snowflake from accessing the files. Incorrect event notification settings (B) would stop events from being sent to the queue/topic. Mismatched file format will not generate any events, and will cause files to be skipped. Incorrect message retention period may result in loss of messages before snowpipe could process them, hence can be the root cause. (C) is incorrect as it does not impact picking up of event notifications. Warehouse size mainly affects processing data and not the notification retrieval.

### NEW QUESTION # 72

You are designing a data pipeline that requires applying a complex scoring algorithm to customer data in Snowflake. This algorithm involves multiple steps, including feature engineering, model loading, and prediction. You want to encapsulate this logic within a reusable component and apply it to incoming data streams efficiently. Which of the following approaches is most suitable and scalable for implementing this scoring logic as a UDF/UDTF, considering real-time data processing and low latency requirements?

- A. A Python UDF that loads a pre-trained machine learning model (e.g., using scikit-learn) and performs predictions on the input data.
- B. A SQL UDF containing a series of nested CASE statements to implement the entire scoring algorithm.
- C. A JavaScript UDF that uses basic JavaScript functions to perform the entire scoring algorithm without external dependencies.
- D. A Java UDTF that leverages a custom Java library for feature engineering and model prediction, deployed as a JAR file to Snowflake's internal stage.
- E. A Python UDTF using Snowpark, leveraging external libraries like 'torch' for accelerated calculations and ML model inference by GPU.

**Answer: E**

Explanation:

Python UDTFs in Snowpark provide a powerful and scalable way to perform complex scoring algorithms, especially when combined with the GPU. Snowpark optimizes data processing within Snowflake's engine, and integration with Anaconda allows for using machine learning libraries such as scikit-learn or Pytorch for model loading and prediction and external libraries like 'torch' for accelerated calculations and ML model inference by GPU. SQL UDFs are not suitable for complex algorithms. JavaScript UDFs lack the necessary functionality and performance for advanced scoring. While Java UDTFs can be used, managing JAR files and potentially less efficient integration can be disadvantages. Using Python with SNOWPARK for GPU is suitable for real-time scoring and low latency.

### NEW QUESTION # 73

You are designing a CI/CD pipeline for your Snowflake data transformations. One stage involves testing a new stored procedure that modifies several tables in your data warehouse. To ensure data integrity and proper rollback capabilities during testing in your development environment, you want to use a combination of cloning and Time Travel. Select the option that represents the most robust strategy for testing with the ability to revert to the original state in case of failures. Choose all that apply.

- A. Before executing the stored procedure in development, create clones of all affected tables and store the clone names in a configuration file for rollback if needed.
- B. Leverage Snowflake's Time Travel feature to create a named snapshot of each affected table immediately before running the stored procedure. Use the 'AT(TIMESTAMP ...)' or 'BEFORE(STATEMENT => ...)' syntax within your testing framework, for easy reversion.
- C. Create a clone of the schema containing all the tables affected by the stored procedure before executing the procedure in the development environment. This will create point in time backup to rollback at schema level.
- D. Create a new database by cloning the existing database before the transformations. Use the cloned database for all development and testing.
- E. Immediately after executing the stored procedure, execute 'UNDROP TABLE' on all the affected tables. If there were errors, it would restore data from the point of the procedure execution.

**Answer: B,C**

Explanation:

Options B and C provide the most robust and Snowflake-native rollback capabilities. Cloning the schema provides a complete and consistent snapshot of all affected tables. Time Travel offers fine-grained control for reverting individual tables to a specific point in time. Option A is less efficient as you would need to manage numerous clones and track them manually which would impact the development. Option D 'UNDROP TABLE' command is for objects that have been dropped already, it is not relevant here. Cloning database in Option E is too resource intensive for regular testing within a CI/CD pipeline.

### NEW QUESTION # 74

A Snowflake table 'CUSTOMER ORDERS' is clustered by 'ORDER DATE'. You have observed the clustering depth increasing over time, impacting query performance. To improve performance, you decide to recluster the table. However, you need to minimize the impact on concurrent DML operations and cost. Which of the following strategies would be MOST effective in managing this recluster process?

- A. Implement a continuous recluster process using Snowpipe to automatically recluster new data as it arrives.
- B. Use 'CREATE OR REPLACE TABLE with SELECT FROM CUSTOMER ORDERS' to rebuild the table with optimized clustering.
- C. Create a new table clustered by 'ORDER\_DATE', copy data in parallel, and then swap tables.
- D. Leverage Snowflake's automatic recluster feature, monitor its performance, and adjust warehouse size as needed.
- E. Recluster the entire table in a single transaction during off-peak hours.

**Answer: D**

Explanation:

Snowflake's automatic recluster feature is designed specifically to address this scenario. It automatically reclusters data in the background, minimizing the impact on concurrent DML operations. Options A and E will lock the table and can cause performance impacts during peak hours. Option B is more complex and requires significant downtime during the swap. Option C describes Snowpipe data ingestion, which isn't recluster. While creating clustered tables with Snowpipe is possible, this option does not address the already unclustered data. Adjusting the warehouse size is relevant to recluster performance but is secondary to enabling the automatic recluster itself.

### NEW QUESTION # 75

You have implemented a row access policy on a 'products' table to restrict access based on the user's group. The policy uses a mapping table 'user\_groups' to determine which products a user is allowed to see. After implementing the policy, users are reporting significant performance degradation when querying the 'products' table. What are the MOST likely causes of this performance issue, and what steps can you take to mitigate them? Select all that apply.

- A. The 'user\_groups' table is not properly indexed, causing slow lookups during policy evaluation. Create an index on the

'username' and 'group' columns of the 'user\_groups' table.

- B. The users do not have sufficient privileges to access the 'user\_groups' table. Grant the necessary SELECT privileges to the users on the 'user\_groupS' table.
- C. The row access policy is interfering with Snowflake's data pruning capabilities. Ensure that the policy expression can be evaluated efficiently by Snowflake's query optimizer by using the 'USING' clause of the ROW ACCESS POLICY.
- D. The row access policy is overly complex and contains computationally expensive functions. Simplify the policy logic and avoid using UDFs or complex subqueries within the policy definition.
- E. The row access policy is causing full table scans on the 'products' table. Review the query patterns and consider adding clustering keys to the 'products' table to improve data access patterns.

**Answer: A,C,D,E**

Explanation:

All options except D are likely causes of performance degradation. A poorly indexed 'user\_groups' table (A) will slow down policy evaluation. Complex policy logic (B) can also impact performance. Interference with data pruning (C) is a common issue with row access policies. Full table scans (E) can be exacerbated by the policy if data is not clustered appropriately. Users needing explicit privileges to 'user\_groups' is not needed since the policy handles that; also using a secure view handles that as well.

## NEW QUESTION # 76

.....

All we want you to know is that people are at the heart of our manufacturing philosophy, for that reason, we place our priority on intuitive functionality that makes our DEA-C02 Exam Question to be more advanced. So with our DEA-C02 guide torrents, you are able to pass the exam more easily in the most efficient and productive way and learn how to study with dedication and enthusiasm, which can be a valuable asset in your whole life. It must be your best tool to pass your exam and achieve your target.

**Exam Sample DEA-C02 Questions:** <https://www.validvce.com/DEA-C02-exam-collection.html>

EXIN DEA-C02 Materials - So you have no reason not to choose it, More importantly, we can assure you that if you use our DEA-C02 certification guide, you will never miss any important and newest information, Snowflake DEA-C02 New Practice Questions Besides, the answers together with questions are authorized and can ensure you pass with ease, Once you become our customers we provide you one-year service warranty, if you fail exam with our DEA-C02 test braindumps please provide the failure score scanned we will arrange refund soon.

Why LinkedIn Is Relevant to Your Business, Discovering the Apple Music Service, EXIN DEA-C02 Materials - So you have no reason not to choose it, More importantly, we can assure you that if you use our DEA-C02 Certification guide, you will never miss any important and newest information.

## ValidVCE Snowflake DEA-C02 Exam Questions Formats

Besides, the answers together with questions are DEA-C02 authorized and can ensure you pass with ease, Once you become our customers we provide you one-year service warranty, if you fail exam with our DEA-C02 test braindumps please provide the failure score scanned we will arrange refund soon.

There will be one version right for you and help you quickly pass the DEA-C02 with ease, so that you can obtain the most authoritative international recognition on your IT ability.

- Free PDF Snowflake - Efficient DEA-C02 New Practice Questions  Download ➤ DEA-C02  for free by simply entering ➡ [www.prepawaypdf.com](http://www.prepawaypdf.com)  website  DEA-C02 Guaranteed Passing
- Free PDF Snowflake - Efficient DEA-C02 New Practice Questions  Search for ( DEA-C02 ) and obtain a free download on **【 [www.pdfvce.com](http://www.pdfvce.com) 】**  DEA-C02 Lead2pass
- Pass Guaranteed Quiz Snowflake - Pass-Sure DEA-C02 - SnowPro Advanced: Data Engineer (DEA-C02) New Practice Questions   [www.exam4labs.com](http://www.exam4labs.com)  is best website to obtain ✓ DEA-C02  ✓  for free download  DEA-C02 Valid Test Bootcamp
- Valid DEA-C02 Test Dumps  DEA-C02 Valid Exam Voucher  DEA-C02 Lead2pass  Search for ➡ DEA-C02  and obtain a free download on ➡ [www.pdfvce.com](http://www.pdfvce.com)   Latest DEA-C02 Dumps Sheet
- Free PDF Snowflake Marvelous DEA-C02 New Practice Questions  Search for  DEA-C02  and download it for free on [ [www.practicevce.com](http://www.practicevce.com) ] website  Latest DEA-C02 Dumps Sheet
- Trustworthy DEA-C02 Exam Content  DEA-C02 Practice Engine  DEA-C02 Lead2pass  Search for ➡ DEA-C02  and download it for free immediately on ⇒ [www.pdfvce.com](http://www.pdfvce.com) ⇐  DEA-C02 Valid Exam Voucher

