

# 100% Pass Quiz Snowflake - Updated NAS-C01 - SnowPro Specialty - Native Apps Training Tools



Compared with the education products of the same type, some users only for college students, some only provide for the use of employees, these limitations to some extent, the product covers group, while our NAS-C01 research material absorbed the lesson, it can satisfy the different study period of different cultural levels of the needs of the audience. For example, if you are a college student, you can study and use online resources through the student column of our NAS-C01 Study Materials, and you can choose to study in your spare time.

Do you long to get the NAS-C01 certification to improve your life? Are you worried about how to choose the NAS-C01 learning product that is suitable for you? If your answer is yes, we are willing to tell you that you are a lucky dog, because you meet us, it is very easy for us to help you solve your problem. The NAS-C01 latest question from our company can help people get their NAS-C01 certification in a short time.

>> NAS-C01 Training Tools <<

## NAS-C01 Real Sheets & Test NAS-C01 Simulator

By using PrepAwayPDF NAS-C01 questions pdf, you will be able to understand the real exam NAS-C01 scenario. It will help you get verified NAS-C01 answers and you will be able to judge your NAS-C01 preparation level for the NAS-C01 exam. More importantly, it will help you understand the real SnowPro Specialty - Native Apps exam feel. You will be able to check the real exam scenario by using this specific NAS-C01 Exam PDF questions. Our NAS-C01 experts are continuously working on including new NAS-C01 questions material and we provide a guarantee that you will be able to pass the NAS-C01 exam on the first attempt.

## Snowflake SnowPro Specialty - Native Apps Sample Questions (Q199-Q204):

### NEW QUESTION # 199

You are developing a Snowflake Native Application that leverages Snowflake's Snowpark API for data processing. The application performs a series of complex transformations on a DataFrame. You need to optimize the application's performance and minimize resource consumption. Which of the following strategies would be MOST effective in achieving this goal?

- A. Employ lazy evaluation by chaining together transformations on the DataFrame and only triggering execution when the final result is needed.
- B. Explicitly specify the execution order of the transformations using 'cache()' on each intermediate DataFrame to force immediate materialization.
- C. Break down the DataFrame processing into smaller steps and store the intermediate results in temporary tables within the application's sandbox database.
- D. Use the method frequently to materialize intermediate DataFrames and inspect the data.
- E. Rely on Snowpark's query optimizer to automatically optimize the entire data processing pipeline without any manual intervention.

Answer: A

Explanation:

Lazy evaluation (C) is the most effective strategy. By chaining transformations and only executing when the final result is needed, Snowpark's query optimizer can optimize the entire pipeline, reducing resource consumption. 'collect()' (A) materializes DataFrames unnecessarily, hindering optimization. While Snowpark does optimize queries (B), relying solely on it is insufficient. Explicit materialization (D) defeats the purpose of lazy evaluation. Storing intermediate results in temporary tables (E) adds overhead and complexity.

#### NEW QUESTION # 200

You are developing a Snowflake Native Application that utilizes Streamlit for its user interface. The application needs to access data from a protected table within the provider account. To enhance security, you want to grant specific privileges to the Streamlit application without exposing the underlying table directly. Which of the following steps are necessary to achieve this, ensuring the principle of least privilege?

- A. Create a secure view on the protected table and grant 'SELECT' privilege on the secure view to the Streamlit application's role.
- B. Use an external function (UDF) owned by the application to query the protected table and grant 'USAGE' privilege on the UDF to the Streamlit application's role.
- C. Grant 'SELECT' privilege directly on the protected table to the Streamlit application's role.
- D. Create a stored procedure with 'EXECUTE AS CALLER' which queries the protected table and grant 'USAGE' privilege on the procedure to the Streamlit application's role.
- E. Create an API integration, implement a service (e.g., AWS Lambda) that queries the data and grant 'USAGE' on the API Integration and appropriate permissions on the service to the Streamlit application's role.

Answer: A,B,D

Explanation:

Granting 'SELECT' directly to the Streamlit app role is against the principle of least privilege (A). Using a secure view (B), a UDF (C) or stored procedure (E) allows you to control precisely what data is exposed and under what conditions. API Integration will work. The service function (Lambda, etc) would then require permissions to query Snowflake. EXECUTE AS CALLER allows the procedure to run with the rights of the caller.

#### NEW QUESTION # 201

Consider the following setup script snippet for a Snowflake Native Application:

```
CREATE SCHEMA IF NOT EXISTS app_data;
CREATE TABLE IF NOT EXISTS app_data.configuration (
  param_name VARCHAR,
  param_value VARCHAR
);
INSERT INTO app_data.configuration (param_name, param_value)
SELECT 'api_endpoint', 'https://default.api.example.com'
WHERE NOT EXISTS (SELECT 1 FROM app_data.configuration WHERE param_name = 'api_endpoint');
```

What is the primary purpose of the 'WHERE NOT EXISTS' clause in the 'INSERT' statement within this setup script?

- A. To ensure that the 'INSERT' statement only executes if the table is empty.
- B. To handle potential errors that may occur if the 'api\_endpoint' parameter is already defined in the consumer's account.
- C. To guarantee that the 'app\_data.configuration' table always contains the latest version of the 'api\_endpoint' parameter, overwriting any existing value.
- D. To prevent duplicate rows with the same 'param\_name' from being inserted into the 'app\_data.configuration' table during subsequent installations or upgrades of the application.
- E. To optimize the performance of the 'INSERT' statement by avoiding unnecessary insertions when the configuration parameter already exists.

Answer: D

Explanation:

Option B is correct. The 'WHERE NOT EXISTS' clause prevents duplicate entries for the 'api\_endpoint' parameter. The setup script might run multiple times during application installations or upgrades, and without this clause, it would insert duplicate rows. Option A is incorrect, as the clause doesn't check if the entire table is empty. Option C is true to some extent, performance is one of the reasons but not the primary reason. Option D, The consumer account has no role to play here. Option E is wrong because it will

do opposite as intended.

### NEW QUESTION # 202

A Snowflake Native Application is designed to process customer data using a series of stored procedures. One of the stored procedures, 'process\_data', requires access to a stage named 'customer\_stage' in the consumer's account to load data. The application role 'app\_role' is intended to manage all access within the application. However, the procedure fails with a permission error during testing in the consumer environment. Which of the following combination of steps is required to fix the problem?

- A. Grant 'READ and ' WRITE privilege on 'customer\_stage' directly to the provider account.
- B. Grant 'USAGE privilege on the database and schema containing 'customer\_stage' to the 'app\_role', and create a new role in the consumer account, grant 'READ on the 'customer\_stage' to this new role, then grant this new role to 'app\_role' .
- C. The stored procedure needs to be defined with 'EXECUTE AS CALLER rights to assume the privileges of the role executing the procedure.
- D. Grant 'USAGE privilege on the database and schema containing 'customer\_stage' to the 'app\_role', and grant 'READ privilege on 'customer\_stage' to the app\_role' .
- E. Grant 'USAGE privilege on the database and schema containing 'customer\_stage' to the provider account, and grant privilege on 'customer\_stage' to the 'app\_role' .

**Answer: C,D**

Explanation:

Options A and E are correct. Option A: The 'app\_role' needs 'USAGE privileges on the database and schema to access objects within them. It also needs 'READ' privilege on the stage to load data from it. Option E: If the stored procedure is accessing objects with the 'OWNER rights instead of 'CALLER rights, then it doesn't matter what the 'app\_role' has privileges to. The procedure needs to be defined to execute as the caller to take advantage of the consumers permissions to external stages. Options B and C are incorrect because granting privileges directly to the provider account or the consumer account is the incorrect pattern for Snowflake Native Apps. The privileges should be managed within the consumer account, usually through an 'app\_role' . Option D includes an unnecessary intermediary role that is not needed.

### NEW QUESTION # 203

You are developing a Snowflake Native Application that provides data enrichment services. A consumer reports that a specific query against a view, 'ENRICHED DATA, occasionally returns incorrect results. You suspect a race condition in the underlying stored procedure that populates a temporary table used by the view. Which of the following strategies would be MOST effective in diagnosing and resolving this issue within the application ?

- A. Replace the temporary table with a permanent table with appropriate locking mechanisms to prevent concurrent modifications. Consumers will need to be notified of a minor version update.
- B. Isolate the portion of the stored procedure that writes to the temporary table into a separate stored procedure with 'EXECUTE AS CALLER privileges.
- C. Implement a retry mechanism with exponential backoff in the stored procedure to handle potential concurrency issues when writing to the temporary table.
- D. Implement logging within the stored procedure to capture the execution flow and intermediate data values. Use 'SYSTEM\$GET PREDECESSORS to track dependencies. Analyze the logs to identify the race condition. Use a session table rather than a temp table.
- E. Use Snowflake's 'GET\_DDL' function to retrieve the SQL definition of the consumer's query and analyze it for potential errors.

**Answer: C,D**

Explanation:

Options A and D are correct. A retry mechanism can mitigate concurrency issues. Logging helps diagnose the race condition. Option B introduces a permanent table, which changes the app's architecture and might have unintended side effects. Option C examines the consumer's query, but the issue is within the provider's application. Option E is unlikely to resolve the race condition and can introduce security concerns. A session table ensures isolation and mitigates concurrency problems.

### NEW QUESTION # 204

.....

