

시험대비SPS-C01최신덤프문제최신버전덤프샘플문제 다운



참고: Itexamdump에서 Google Drive로 공유하는 무료, 최신 SPS-C01 시험 문제집이 있습니다:
<https://drive.google.com/open?id=1RvLTP5xAouAmuLNgPEJ7zOU6DagAC6IZ>

우리는 여러분이 시험패스는 물론 또 일년무료 업데이트서비스를 제공합니다.만약 시험에서 실패했다면 우리는 덤프비용전액 환불을 약속 드립니다.하지만 이런 일은 없을 것입니다.우리는 우리덤프로 100%시험패스에 자신이 있습니다. 여러분은 먼저 우리 Itexamdump사이트에서 제공되는 Snowflake인증SPS-C01시험덤프의 일부인 데모 즉 문제와 답을 다운받으셔서 체험해보실 수 있습니다.

다른 사이트에서도 Snowflake SPS-C01인증 시험관련 자료를 보셨다고 믿습니다.하지만 우리 Itexamdump의 자료는 차원이 다른 완벽한 자료입니다.100%통과 율은 물론Itexamdump을 선택으로 여러분의 직장생활에 더 나은 개편을 가져다 드리며 ,또한Itexamdump를 선택으로 여러분은 이미 충분한 시험준비를 하였습니다.우리는 여러분이 한번에 통과하게 도와주고 또 일년무료 업데이트서비스도 드립니다.

>> SPS-C01최신 덤프문제 <<

인기자격증 SPS-C01최신 덤프문제 시험대비 공부자료

Itexamdump에서 출시한 Snowflake인증 SPS-C01덤프는 실제시험문제 커버율이 높아 시험패스율이 가장 높습니다. Snowflake인증 SPS-C01시험을 통과하여 자격증을 취득하면 여러방면에서 도움이 됩니다. Itexamdump에서 출시한

Snowflake인증 SPS-C01덤프를 구매하여Snowflake인증 SPS-C01시험을 완벽하게 준비하지 않으실래요? Itexamdump의 실력을 증명해드릴게요.

최신 Snowflake Certification SPS-C01 무료 샘플문제 (Q269-Q274):

질문 # 269

You have a Snowpark Python application that reads data from a Snowflake table, performs a complex transformation using a User-Defined Table Function (UDTF), and then writes the transformed data back to a new Snowflake table. The UDTF is defined as follows:

```
from snowflake.snowpark.functions import lit
```

```
class MyUDTF:
    def __init__(self, factor: int):
        self.factor = factor

    def process(self, row: Row) -> Iterable[Tuple[int, str]]:
        value = row[0] * self.factor
        yield (value, f"Value multiplied by {self.factor}")

    def end_partition(self):
        yield (None, "END")
```

```
my_udtf = udf(MyUDTF, return_type=T.StructType([T.StructField("RESULT", T.IntegerType()), T.StructField("MESSAGE", T.StringType())]), input_types=[T.IntegerType()])
```

You need to optimize the performance of this application. Which of the following strategies would be MOST effective in reducing the execution time of the UDTF?

- A. Reduce the number of rows in the input table by applying a filter before calling the UDTF.
- B. Cache the input DataFrame before applying the UDTF using 'df.cache()'.
- C. Use a vectorized UDTF instead of a standard UDTF, processing data in batches. (Assume 'vectorized=True' is a valid parameter for 'udf'.)
- D. Replace the UDTF with a standard User-Defined Function (UDF) as UDFs are inherently faster.
- E. Increase the warehouse size to the largest available option before running the Snowpark application.

정답: C

설명:

Vectorized UDTFs process data in batches, which can significantly improve performance compared to processing each row individually. This is especially true for complex transformations. Increasing warehouse size (A) can help but might not be as efficient as vectorization. Reducing input data (B) is always a good practice, but vectorization provides a more direct performance boost to the UDTF execution. Standard UDFs (D) are not generally faster than UDTFs, especially when dealing with table transformations. Caching (E) can help if the DataFrame is reused multiple times, but it doesn't directly optimize the UDTF's performance.

질문 # 270

You have developed a Snowpark Python application that needs to connect to an external REST API to enrich data during a transformation. The API requires authentication using an API key stored securely. Which of the following approaches is the MOST secure and recommended way to manage the API key within the Snowpark environment?

- A. Store the API key in a Snowflake Secret Object and retrieve it within the Snowpark Python code using the function.
- B. Store the API key as an environment variable within the Snowflake session.
- C. Hardcode the API key directly into the Snowpark Python code.
- D. Encrypt the API key using a third-party encryption library and store it in a Snowflake table.
- E. Store the API key in a secure vault outside of Snowflake and retrieve it using a custom Snowflake external function.

정답: A

설명:

Option C is the most secure and recommended approach. Snowflake Secret Objects provide a secure way to store and manage sensitive information like API keys. The function allows you to retrieve the key within your Snowpark code without exposing it directly. Option A is highly insecure. Option B is less secure than using Secret Objects, as environment variables can be accessed more easily. Option D adds complexity and doesn't provide the same level of security as Secret Objects. Option E introduces external dependencies and requires managing another system, making it less desirable than using built-in Snowflake features.

질문 # 271

You have a Snowflake table 'PRODUCT CATALOG' with columns 'PRODUCT ID', 'PRODUCT NAME', and 'CATEGORY ID'. You also have a table 'CATEGORY' with 'CATEGORY ID' and 'CATEGORY NAME'. You need to create a Snowpark

DataFrame that joins these two tables and includes only 'PRODUCT NAME' and 'CATEGORY NAME'. Assume a Snowpark Session object named 'session' is available. Which code snippet demonstrates creating the DataFrame using Snowpark's join functionality and column selection while using the 'table' method?

- A.

```
product_df = session.table('PRODUCT_CATALOG')
category_df = session.table('CATEGORY')
joined_df = product_df.join(category_df, product_df['CATEGORY_ID'] == category_df['CATEGORY_ID'])
result_df = joined_df.select(col('PRODUCT_NAME'), col('CATEGORY_NAME'))
```

- B.

```
product_df = session.table('PRODUCT_CATALOG')
category_df = session.table('CATEGORY')
joined_df = product_df.join(category_df, product_df['CATEGORY_ID'] == category_df['CATEGORY_ID'])
result_df = joined_df.select('PRODUCT_NAME', 'CATEGORY_NAME')
```

- C.

```
product_df = session.table('PRODUCT_CATALOG')
category_df = session.table('CATEGORY')
joined_df = product_df.join(category_df, product_df.col('CATEGORY_ID') == category_df.col('CATEGORY_ID'))
result_df = joined_df.select('PRODUCT_NAME', 'CATEGORY_NAME')
```

```
product_df = session.table('PRODUCT_CATALOG')
category_df = session.table('CATEGORY')
joined_df = product_df.join(category_df, 'CATEGORY_ID')
result_df = joined_df[['PRODUCT_NAME', 'CATEGORY_NAME']]
```

- D.

```
product_df = session.table('PRODUCT_CATALOG')
category_df = session.table('CATEGORY')
joined_df = product_df.join(category_df, product_df.PRODUCT_ID == category_df.CATEGORY_ID)
result_df = joined_df.select(product_df.PRODUCT_NAME, category_df.CATEGORY_NAME)
```

- E.

정답: A

설명:

Option D correctly joins the tables using the 'join' method with a column expression specifying the join condition. It also correctly selects only the desired columns using and option A incorrectly uses 'PRODUCT_ID' instead of for the join and needs 'col()' to reference columns. option B selects 'PRODUCT_NAME', 'CATEGORY_NAME' without using 'col()' which will cause an error since the join brings duplicate column names from product_df and category_df. Option C uses Pythonic '['PRODUCT_NAME', 'CATEGORY_NAME']' which only works on Pandas dataframes. option E attempts to use col() to reference the columns in the join condition which is incorrect as it is a column expression needs to use fully qualified table name for the same named column in two dataframes; moreover selection is missing col().

질문 # 272

You are tasked with building a Snowpark function to perform an upsert operation on a Snowflake table using a DataFrame. The function should take the target table name, a staging DataFrame, a join key column, and a list of columns to update. The function needs to handle potential schema evolution (i.e., columns may be added or removed from either the target table or the staging DataFrame) gracefully without causing the entire upsert to fail. Which of the following approaches, or combinations of approaches, would best address this requirement?

- A. Use the 'exceptAll' to ensure that there are no schema evolution issues.
- B. Dynamically generate the SQL 'MERGE' statement within the function, comparing the columns present in the target table and the staging DataFrame, and only including those columns that exist in both.
- C. Before the 'merge' operation, use 'DataFrame.select' on the staging DataFrame to project only the columns that exist in the target table.
- D. Rely on Snowflake's automatic schema detection during the 'merge' operation to automatically adapt to schema changes.
- E. Before the merge, create a temporary table with the exact schema of the target table, insert all the data from the DataFrame into it, and then use the temporary table as source for the merge. Handle the schema evolution with dynamic sql if required.

정답: B,C

설명:

Approaches A and D are the most suitable for handling schema evolution during an upsert operation. Approach A involves dynamically generating the SQL MERGE statement by inspecting the schemas of both the target table and the staging DataFrame. This ensures that only the common columns are included in the update and insert clauses, preventing errors due to missing columns. Approach D suggests projecting the staging DataFrame to only include the columns that exist in the target table using DataFrame.select'. This effectively harmonizes the schema of the staging data with the target table's schema, avoiding issues during the 'merge' operation. While Snowflake does have some schema evolution capabilities, explicitly handling it in the code provides more control and predictability.

질문 # 273

You have a Snowpark application processing streaming data from an event table. You observe that the application frequently fails with transient errors related to network connectivity or Snowflake service unavailability. You want to implement a robust error handling strategy to ensure the application can recover from these transient failures without losing data. Which of the following approaches would be MOST appropriate and effective in this scenario, ensuring idempotent processing?

- A. Utilize Snowpark's 'cache()' method to cache the intermediate DataFrame results in memory, reducing the impact of transient failures.
- B. Use Snowflake's built-in retry mechanism for SQL queries by setting the 'CLIENT_SESSION_PARAMETER to a non-zero value.
- C. Implement exponential backoff and jitter in your retry logic when catching exceptions during Snowpark operations. Store the last successfully processed event ID in a metadata table and resume processing from that point after a retry. Ensure all operations are idempotent.
- D. Implement a message queue (e.g., Kafka, SQS) to buffer the incoming event data. The Snowpark application consumes data from the queue, allowing for retries and ensuring no data is lost during transient failures.
- E. Implement a try-except block around the Snowpark DataFrame operations, logging the error and retrying the entire application from the beginning upon failure.

정답: C,D

설명:

Implementing a message queue provides a buffer that isolates the Snowpark application from transient data source failures. E is correct because adding an exponential backoff mechanism with jitter is crucial to prevent overwhelming the system with retries and helps to ensure idempotent processing. Option B can address some internal Snowflake errors, but not connectivity issues. The other approaches do not address data loss or idempotent operation.

질문 # 274

.....

Snowflake인증SPS-C01시험을 패스하여 자격증을 취득한다면 여러분의 미래에 많은 도움이 될 것입니다.Snowflake 인증SPS-C01시험자격증은 기업계에서도 아주 인지도가 높고 또한 알아주는 시험이며 자격증 하나로도 취직은 문제없다고 볼만큼 가치가 있는 자격증 이죠.Snowflake인증SPS-C01시험은 여러분이 지식테스트시험입니다.

SPS-C01퍼펙트 최신 덤프모음집: <https://www.itexamdump.com/SPS-C01.html>

우리Itexamdump에서는 빠른 시일 내에Snowflake SPS-C01관련 자료를 제공할 수 있습니다, Itexamdump SPS-C01퍼펙트 최신 덤프모음집 덤프의 문제와 답은 모두 제일 정확합니다, Snowflake SPS-C01최신 덤프문제 ITExamDump덤프는 IT전문가들이 최신 실러버스에 따라 몇년간의 노하우와 경험을 충분히 활용하여 연구제작해낸 시험대비자료입니다, IT인증시험을 Snowflake Certified SnowPro Specialty - Snowpark덤프로 준비해야만 하는 이유는 SPS-C01덤프는 IT업계 전문가들이 실제 SPS-C01시험문제를 연구하여 시험문제에 대비하여 예상문제를 제작했다는 점에 있습니다, 우리 Itexamdump SPS-C01퍼펙트 최신 덤프모음집을 선택해주신다면 우리는 최선을 다하여 여러분이 꼭 한번에 시험을 패스할 수 있도록 도와드리겠습니다.만약 여러분이 우리의 인증시험덤프를 보시고 시험이랑 틀려서 패스를 하지 못하였다면 우리는 무조건 덤프비용전부를 환불해드립니다.

그럼 끝이 없잖아요, 안 그래도 가기 전에 한 번은 더 보고 싶었던 참이었는데 잘된 일이었다, 우리Itexamdump에서는 빠른 시일 내에Snowflake SPS-C01관련 자료를 제공할 수 있습니다, Itexamdump 덤프의 문제와 답은 모두 제일 정확합니다.

시험패스 가능한 SPS-C01최신 덤프문제 최신버전 덤프샘플문제 다운 받

