Pass Guaranteed 2025 Snowflake Efficient DEA-C02: SnowPro Advanced: Data Engineer (DEA-C02) Braindumps

Our online test engine and the windows software of the DEA-C02 guide materials can evaluate your exercises of the virtual exam and practice exam intelligently. Our calculation system of the DEA-C02 study engine is designed subtly. Our evaluation process is absolutely correct. We are strictly in accordance with the detailed grading rules of the real exam. And our pass rate of the DEA-C02 Exam Questions are high as 98% to 100%, it is unique in the market.

Are you planning to crack the Snowflake SnowPro Advanced: Data Engineer (DEA-C02) DEA-C02 certification test in a short time and don't know how to prepare for it? ITCertMagic has updated DEA-C02 Dumps questions for the applicants who want to prepare for the Snowflake DEA-C02 Certification test successfully within a few days. This study material is available in three different formats that you can trust to crack the Snowflake DEA-C02 certification test with ease.

>> DEA-C02 Braindumps <<

Test Snowflake DEA-C02 Simulator Fee | Latest DEA-C02 Exam Cost

Customizable SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) exam conditions in such a way that you can create your desired DEA-C02 exam with pre-determined questions and exam duration. You will be able to see instant results after going through the DEA-C02 Practice Exam To confirm the product licence. For customer satisfaction, ITCertMagic has also designed a SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) demo version so the candidate can assure the reliability of the Snowflake PDF Dumps.

Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Sample Questions (Q63-Q68):

NEW QUESTION # 63
You are developing a Snowpark Python stored procedure that performs complex data transformations on a large dataset stored in a Snowflake table named 'RAW SALES'. The procedure needs to efficiently handle data skew and leverage Snowflake's distributed processing capabilities. You have the following code snippet:

```
import snowflake.snowpark as snowpark
from snowflake.snowpark.functions import col, sum

def transform_sales(session: snowpark.Session) -> str:
    df = session.table("RAW_SALES")
    # Assume RAW_SALES has columns: sale_id, product_id, sale_date, sale_amoun

    # Your transformation logic here

    df.write.mode("overwrite").save_as_table("TRANSFORMED_SALES")
    return "Transformation completed successfully"
```


Which of the following strategies would be MOST effective to optimize the performance of this Snowpark stored procedure, specifically addressing potential data skew in the 'product id' column, assuming 'product_id' is known to cause uneven data distribution across Snowflake's micro-partitions?

- A. Increase the warehouse size significantly to compensate for the data skew and improve overall processing speed without modifying the partitioning strategy.
- B. Combine salting with repartitioning by adding a random number to the 'product_id' before repartitioning, then removing the salt after the transformation to break up the skew. Then, enable automatic clustering on the 'TRANSFORMED SALES' table.
- C. Utilize Snowflake's automatic clustering on the 'TRANSFORMED_SALES table by specifying 'CLUSTER BY when creating or altering the table to ensure future data is efficiently accessed.
- D. Use the 'pandas' API within the Snowpark stored procedure to perform the transformation, as 'pandas' automatically optimizes for data skew.
- E. Implement a custom partitioning strategy using before the transformation logic to redistribute data evenly across the cluster.

Answer: B

Explanation:
Option E is the most effective solution. Salting breaks up data skew before repartitioning. Automatic clustering on the transformed table optimizes future queries. Repartitioning redistributes the data across Snowflake's processing nodes, and Automatic Clustering will help in maintaining performance as the data changes in TRANSFORMED_SALES table over time. Option A, without salting, may still be inefficient due to the initial skew. Option B improves query performance but doesn't address the initial transformation skew. Option C is incorrect because 'pandas' in Snowpark does not automatically handle data skew at the Snowflake level. Option D is a costly workaround that doesn't fundamentally solve the skew problem.

NEW QUESTION # 64
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 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.
- B. 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.
- C. 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.
- D. 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.
- 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: B,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 # 65
A Snowflake data pipeline ingests data from multiple external sources into a RAW DATA table. A transformation process then moves the data to a ANALYTICS DATA table, applying several complex UDFs written in Java and Python for data cleansing and enrichment. Performance is significantly slower than expected. Which combination of techniques would BEST improve the performance of this transformation pipeline?

- A. Use external functions instead of UDFs to offload the processing to an external compute environment and configure auto-scaling for the virtual warehouse.
- B. Increase the virtual warehouse size and re-cluster the ANALYTICS DATA table based on the most frequently filtered columns.
- C. Rewrite the UDFs in SQL or Snowpark Python/Java for better integration with the Snowflake engine and leverage vectorization where possible; cache intermediate results using temporary tables.
- D. Reduce the number of UDF calls by consolidating them into a single, more complex UDF. Replace the transformation pipeline with a series of COPY INTO statements.
- E. Implement data partitioning in the RAW DATA table based on ingestion time and switch to using stored procedures instead of transformation pipelines.

Answer: C

Explanation:
Rewriting UDFs in SQL or Snowpark allows the Snowflake engine to optimize them more effectively. Using temporary tables to cache intermediate results prevents redundant computation. Option A helps, but addressing UDF performance is more crucial. Option C might add complexity. Option D is irrelevant to UDF performance. Option E is counterproductive; consolidating UDFs can reduce parallelism and COPY INTO is for initial data loading, not transformation.

NEW QUESTION # 66
You have a table named 'EMPLOYEES with a retention period of 1 day. You accidentally deleted several important rows from this table, but you need to recover the data'. You know the deletion occurred 25 hours ago. What actions should be taken to attempt to recover the deleted data, and what outcome can you expect? Assume you are working in an Enterprise edition of Snowflake account.

- A. Since its Enterprise edition of Snowflake account, the Time travel and cloning will work with 7 days retention period, hence attemtp clone table using Time Travel and recover data successfully
- B. Attempt to clone the table using Time Travel to a point in time before the deletion, then extract the deleted rows. Expect the recovery to be successful as long as the deletion occurred within the data retention period.
- C. Attempt to use Time Travel or cloning to recover the data. Expect the recovery to fail because the deletion occurred outside the I-day data retention period.
- D. Attempt to use UNDROP TABLE command if the table was dropped. Expect the recovery to be successful as long as the deletion occurred within the data retention period.
- E. Attempt to use Time Travel to query the table before the deletion and re-insert the deleted rows. Expect the recovery to be successful as long as the deletion occurred within the data retention period.

Answer: C

Explanation:
Option D is the correct answer. Snowflake's Time Travel feature allows data recovery within the defined data retention period. Since the deletion occurred 25 hours ago and the retention period is only 1 day (24 hours), the data is no longer recoverable using Time Travel or cloning. It's important to note that even if the deletion was only 25 hours ago, it still outside of the I-day data retention period.

NEW QUESTION # 67
You are tasked with creating a development environment from a production database in Snowflake. The production database is named 'PROD DB' and contains several schemas, including 'CUSTOMER DATA' and 'PRODUCT DATA'. You want to create a clone of the 'PROD DB' database named 'DEV DB', but you only need the 'CUSTOMER DATA' schema for development purposes and all the data should be masked with a custom UDF 'MASK EMAIL' for 'email' column in 'CUSTOMER' table. The 'email' column is VARCHAR. Which of the following sequences of SOL statements would achieve this in Snowflake? Note: UDF MASK EMAIL already exists in the account.

- A.

```
CREATE DATABASE DEV_DB CLONE PROD_DB;
DROP SCHEMA DEV_DB.PRODUCT_DATA;
CREATE OR REPLACE TABLE DEV_DB.CUSTOMER_DATA.CUSTOMER AS
SELECT  ,
MASK_EMAIL(email) as email
FROM PROD_DB.CUSTOMER_DATA.CUSTOMER;
```

- B.

```
CREATE DATABASE DEV_DB CLONE PROD_DB;
DROP SCHEMA DEV_DB.PRODUCT_DATA;
CREATE OR REPLACE TABLE DEV_DB.CUSTOMER_DATA.CUSTOMER CLONE PROD_DB.CUSTOMER_DATA.CUSTOMER;
ALTER TABLE DEV_DB.CUSTOMER_DATA.CUSTOMER
MODIFY COLUMN email SET MASKING POLICY mask_email;
```

- C.

```
CREATE DATABASE DEV_DB CLONE PROD_DB;
DROP SCHEMA DEV_DB.PRODUCT_DATA;
CREATE OR REPLACE TABLE DEV_DB.CUSTOMER_DATA.CUSTOMER AS
SELECT   FROM PROD_DB.CUSTOMER_DATA.CUSTOMER;
UPDATE DEV_DB.CUSTOMER_DATA.CUSTOMER SET email = MASK_EMAIL(email);
```

- D.

```
CREATE DATABASE DEV_DB CLONE PROD_DB;
DROP SCHEMA DEV_DB.PRODUCT_DATA;
CREATE OR REPLACE VIEW DEV_DB.CUSTOMER_DATA.CUSTOMER AS
SELECT  ,
MASK_EMAIL(email) as email
FROM PROD_DB.CUSTOMER_DATA.CUSTOMER;
```

- E.

```
CREATE DATABASE DEV_DB CLONE PROD_DB;
DROP SCHEMA DEV_DB.PRODUCT_DATA;
CREATE OR REPLACE TABLE DEV_DB.CUSTOMER_DATA.CUSTOMER AS
SELECT  ,
MASK_EMAIL(email) as email
FROM PROD_DB.CUSTOMER_DATA.CUSTOMER;
ALTER TABLE DEV_DB.CUSTOMER_DATA.CUSTOMER DROP COLUMN email;
ALTER TABLE DEV_DB.CUSTOMER_DATA.CUSTOMER ADD COLUMN email VARCHAR;
```

Answer: B

Explanation:
Option B is the most appropriate solution. It clones the entire production database, drops the unnecessary schema, then clone table from PROD and after cloning, it uses masking policy on email column on the cloned DEV environment. Option A is incorrect because you cannot use MASK EMAIL while createing the table. Option C requires to drop and add column again, option D, using a view will not permanently mask data at the storage level. And Option E updates the table after cloning, which consumes resources and isn't as elegant as using a masking policy.

NEW QUESTION # 68
......

Through the stimulation of the DEA-C02 real exam the clients can have an understanding of the mastery degrees of our DEA-C02 exam practice question in practice. Thus our clients can understand the abstract concepts in an intuitive way. In the answers, our experts will provide the authorized verification and detailed demonstration so as to let the learners master the latest information timely and follow the trend of the times. All we do is to integrate the most advanced views into our DEA-C02 Test Guide.

Test DEA-C02 Simulator Fee: https://www.itcertmagic.com/Snowflake/real-DEA-C02-exam-prep-dumps.html

The PDF file is downloadable and printable with 90 days of free update options of DEA-C02 exam dumps that show that the students have plenty of opportunities to practice and pass the DEA-C02 Snowflake Certified Specialist certification exam, Snowflake DEA-C02 Braindumps Never worry, I will tell you why, Snowflake DEA-C02 Braindumps You can choose your preferential one based on your interest.

Keep your composure, and act professionally, Perform wireless DEA-C02 reconnaissance, The PDF file is downloadable and printable with 90 days of free update options of DEA-C02 Exam Dumps that show that the students have plenty of opportunities to practice and pass the DEA-C02 Snowflake Certified Specialist certification exam.

Use Snowflake DEA-C02 PDF Questions To Take Exam With Confidence

Never worry, I will tell you why, You can choose your preferential one based on your interest, If you need more information for the Snowflake DEA-C02 Exam, check our page Snowflake DEA-C02 Dumps.

And the DEA-C02 quiz guide's experts still have the ability to master propositional trends.

- DEA-C02 Latest Exam Cram ☎ Braindumps DEA-C02 Torrent □ DEA-C02 Practice Online □ Go to website ▹ www.pass4leader.com ◃ open and search for 【 DEA-C02 】 to download for free □DEA-C02 Practice Online
- Snowflake DEA-C02 Three Formats for Preparations □ Download □ DEA-C02 □ for free by simply searching on ➡ www.pdfvce.com □□□ □DEA-C02 Visual Cert Exam
- Excellent Web-Based Snowflake DEA-C02 Practice Exam □ The page for free download of □ DEA-C02 □ on ➠ www.examcollectionpass.com □ will open immediately □Training DEA-C02 Solutions
- Training DEA-C02 Solutions □ New DEA-C02 Test Format □ 100% DEA-C02 Exam Coverage □ Download " DEA-C02 " for free by simply entering □ www.pdfvce.com □ website □DEA-C02 Exam Sample Questions
- DEA-C02 Exam Sample Questions □ Reliable DEA-C02 Real Test □ DEA-C02 New APP Simulations □ Easily obtain ▶ DEA-C02 ◀ for free download through ➤ www.vceengine.com □ □DEA-C02 Examcollection Questions Answers
- DEA-C02 Interactive Course □ Reliable DEA-C02 Test Vce □ Valid Braindumps DEA-C02 Questions □ Download ➡ DEA-C02 □□□ for free by simply entering ➥ www.pdfvce.com □ website □DEA-C02 Examcollection Questions Answers
- Latest SnowPro Advanced: Data Engineer (DEA-C02) vce dumps - DEA-C02 prep4sure exam □ Go to website ✔ www.testsimulate.com □✔□ open and search for □ DEA-C02 □ to download for free □Valid Braindumps DEA-C02 Questions
- Reliable DEA-C02 Test Bootcamp □ Standard DEA-C02 Answers □ Training DEA-C02 Solutions □ Search for ➠ DEA-C02 □ and download it for free immediately on 《 www.pdfvce.com 》 □DEA-C02 Exam Pattern
- Training DEA-C02 Solutions □ New DEA-C02 Test Format □ DEA-C02 Interactive Course □ ☀ www.pass4leader.com □☀□ is best website to obtain □ DEA-C02 □ for free download □Valid Braindumps DEA-C02 Questions
- Reliable DEA-C02 Real Test □ DEA-C02 Visual Cert Exam □ DEA-C02 New APP Simulations □ Open ▹ www.pdfvce.com ◃ enter ➡ DEA-C02 □□□ and obtain a free download □Reliable DEA-C02 Test Bootcamp
- Fantastic Snowflake DEA-C02 Braindumps - www.getvalidtest.com Free Download □ Easily obtain ➥ DEA-C02 □ for free download through ➥ www.getvalidtest.com □ □Valid Test DEA-C02 Tips
- studytonic.com, www.yuliancaishang.com, www.stes.tyc.edu.tw, lms.ait.edu.za, www.stes.tyc.edu.tw, academy.webrocket.io, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, lms.ait.edu.za, www.stes.tyc.edu.tw, courses.dorisgideon.com, Disposable vapes

What's more, part of that ITCertMagic DEA-C02 dumps now are free: https://drive.google.com/open?id=1w1q5d96cTsIYtZf7DTBPWE1FZKSVz7bh