

Snowflake DAA-C01 Latest Test Vce: SnowPro Advanced: Data Analyst Certification Exam - Prep4sures Helps you Prepare Easily



2025 Latest Prep4sures DAA-C01 PDF Dumps and DAA-C01 Exam Engine Free Share: <https://drive.google.com/open?id=1GPIGJW5UVk07QxjSea7oRJCPM0pHSalm>

If you have interests with our DAA-C01 practice materials, we prefer to tell that we have contacted with many former buyers of our DAA-C01 exam questions and they all talked about the importance of effective DAA-C01 practice material playing a crucial role in your preparation process. Our DAA-C01 practice materials keep exam candidates motivated and efficient with useful content based wholly on the real DAA-C01 guide materials. There are totally three versions of DAA-C01 practice materials which are the most suitable versions for you: pdf, software and app versions.

In the course of your study, the test engine of DAA-C01 actual exam will be convenient to strengthen the weaknesses in the learning process. This can be used as an alternative to the process of sorting out the wrong questions of DAA-C01 learning torrent in peacetime learning, which not only help you save time, but also makes you more focused in the follow-up learning process with our DAA-C01 Learning Materials. Choose our DAA-C01 guide materials and you will be grateful for your right decision.

>> DAA-C01 Latest Test Vce <<

Reliable Snowflake DAA-C01 Test Preparation - Reliable DAA-C01 Test Notes

Now you don't need to spend too much time and money preparing for the Snowflake DAA-C01 test. Just get the latest DAA-C01 exam dumps from Prep4sures and prepare the DAA-C01 test in a very short time. These Customer Experience (Snowflake) DAA-C01 updated dumps will eliminate your risk of failing and enhance your chance of success in the Prep4sures test. Using Snowflake DAA-C01 Exam study material you will gain the best Snowflake DAA-C01 exam knowledge and you will attempt the final DAA-C01 certification test with confidence.

Snowflake SnowPro Advanced: Data Analyst Certification Exam Sample

Questions (Q36-Q41):

NEW QUESTION # 36

When choosing between using a dimensional model and a flattened dataset for BI requirements in Snowflake, what considerations impact the final decision? (Select all that apply)

- A. Data denormalization and redundancy needs
- B. Query performance expectations
- C. User access control requirements
- D. Query execution plan constraints

Answer: A,B

Explanation:

Considerations such as query performance expectations and data denormalization needs influence the choice between a dimensional model and a flattened dataset, affecting query optimization and data structure suitability for BI purposes.

NEW QUESTION # 37

A data analyst observes a sudden and significant drop in sales for a particular product category within a Snowflake database. Initial investigations point to a possible data quality issue. Which of the following steps provides the MOST effective and efficient diagnostic approach using Snowflake features to pinpoint the root cause of the anomaly, focusing on data integrity?

- A. Rebuild the entire product sales table from the raw data source to ensure data consistency, as the drop in sales suggests widespread data corruption. Compare the rebuilt table with the current table.
- B. Analyze query history in Snowflake to identify any recent DML (Data Manipulation Language) operations (e.g., UPDATE, DELETE) performed on the sales table, focusing on operations targeting the affected product category. Correlate these operations with the timeframe of the sales drop.
- C. Disable all ETL processes and re-ingest the sales data overnight. This will overwrite the potential problematic sales data and resolve the drop in sales.
- D. Create a clone of the sales table. Run a full table scan on the clone and perform descriptive statistics on the cloned data to detect statistical outliers without impacting the production environment. Compare with known good statistics.
- E. Utilize Snowflake's Time Travel feature to compare the data before and after the suspected anomaly. Query both datasets and perform set operations like EXCEPT or MINUS to identify changed records within the product category. Analyze those changed records for patterns.

Answer: B,E

Explanation:

Options A and C are the most effective. A utilizes Time Travel for direct data comparison before and after the incident, allowing for focused analysis on changed records. C investigates DML operations, which could directly explain data changes. Option B is inefficient and disruptive. Option D, while helpful, might not pinpoint the cause of a data corruption issue as fast as A and C. Option E is a poor solution, as it doesn't identify the root cause of the issue and it leads to potential data lost from the transactions between the last successful load and when the ETL processes were disabled.

NEW QUESTION # 38

You have Parquet files containing customer order data in an external stage. You need to create a secure view in Snowflake that filters the data based on user roles. Only users with the 'SALES REP' role should see orders from their assigned region, while 'MANAGER' roles should see all data'. Assuming you have set up appropriate role hierarchies, how would you implement this using row-level security policies with Parquet external tables?

- A. Create a secure view that queries the external table. Then, define a row access policy that filters the data based on the user role and the region, and apply the policy to the secure view. The policy should use to determine the appropriate filtering criteria. Delegate the ownership to 'ACCOUNTADMIN'.
- B. Create a secure view that directly queries the Parquet files using to access the region information. Apply a WHERE clause in the view definition using to filter based on the user's role. No RLS policy needed.
- C. Create an external table with a schema defined according to the Parquet files. Define a masking policy that hides the sensitive information for users without proper roles. Create a standard view on top of the external table. Apply the masking policy to a column inside view.
- D. Create a secure view that queries the Parquet files directly using 'JSON EXTRACT PATH TEXT' to access the region

information. Apply a JOIN operation with a role table to filter the data based on the user's role.

- E. Create a standard view querying the external table. Create a row access policy that filters the data based on the user role and the region. Apply the policy to the external table. Delegate the ownership to 'ACCOUNTADMIN'.

Answer: A

Explanation:

Row Access Policies (RAP) are designed for Row Level Security, and they are applied to tables or views. Since the data is in an external table accessed through a view, the policy should be attached to the secure view. The ' function is used to determine the user's current role within the policy's logic. Secure Views protect the underlying data and policy logic from unauthorized access. Delegating ownership to 'ACCOUNTADMIN' is not strictly required but a good security practice. Masking policies aren't the correct choice here as it is used to hide data from the user.

NEW QUESTION # 39

A ride-sharing company wants to analyze the density of ride requests in different city areas. They have a table 'RIDES' with a 'LOCATION' (GEOGRAPHY) column representing the pickup location of each ride. They want to divide the city into a grid of hexagonal cells and count the number of rides originating in each cell. Which sequence of steps would achieve this, assuming the 'city_boundary' is defined and accessible as a GEOGRAPHY object?

- A. 1. Generate a grid of hexagonal GEOGRAPHY objects covering the 'city_boundary' using 'ST_HEXGRID'. 2. Use 'ST_INTERSECTS' to determine which rides intersect with each hexagonal cell. 3. Aggregate the ride counts for each cell.
- B. 1. Generate a grid of hexagonal GEOGRAPHY objects covering the 'city_boundary' using 'ST_SPHERE_GRID'. 2. Use 'ST_CONTAINS' to determine which rides fall within each hexagonal cell. 3. Aggregate the ride counts for each cell.
- C. 1. Generate a grid of rectangular GEOGRAPHY objects covering the 'city_boundary' using 'ST_GRID'. 2. Use 'ST_CONTAINS' to determine which rides fall within each rectangular cell. 3. Aggregate the ride counts for each cell.
- D. 1. Generate a grid of hexagonal GEOGRAPHY objects covering the 'city_boundary' using 'ST_HEXGRID'. 2. Use 'ST_WITHIN' to determine which hexagonal cell each ride falls within. 3. Aggregate the ride counts for each cell.
- E. 1. Generate a grid of hexagonal GEOGRAPHY objects covering the 'city_boundary' using 'ST_HEXGRID'. 2. Use 'ST_CONTAINS' to determine which hexagonal cell contains each ride. 3. Aggregate the ride counts for each cell.

Answer: E

Explanation:

'ST_HEXGRID' is the correct function for generating a hexagonal grid. 'ST_CONTAINS' is used to check if a hexagonal cell contains a ride's location. Aggregating the ride counts for each cell provides the density information.

NEW QUESTION # 40

An organization needs to provide daily summaries of key metrics to its executive team via email. They currently have a complex SQL query that generates the report, but the query takes several minutes to run. The executive team needs the data before 8 AM daily. The data warehouse is also used for other reporting and analytics tasks throughout the day. Given these constraints, what is the MOST efficient and reliable approach to operationalize this report?

- A. Create a Snowpipe to continuously ingest data into a separate table. Write a simplified SQL query on the target table to generate the report. Schedule a task to execute the query and email the results.
- B. Use a third-party ETL tool to extract the data, transform it, and generate the report. Schedule the ETL job to run nightly and email the report.
- C. Use Snowflake's Data Marketplace to find a pre-built reporting solution that matches their requirements.
- D. Create a materialized view based on the complex SQL query. Schedule a task to refresh the materialized view nightly before 8 AM. Create a separate, simpler query against the materialized view to generate the report and email it.
- E. Schedule a Snowflake task to execute the complex SQL query directly and email the results using Snowflake's email integration. Ensure the task runs on a large warehouse.

Answer: D

Explanation:

Creating a materialized view (D) is the most efficient approach. Materialized views pre-compute the results, so the nightly task refreshes the pre-calculated data. This ensures fast report generation before 8 AM, even with a complex underlying query. Running the complex query directly (A) might exceed the time limit. Using a third-party ETL tool (C) adds unnecessary complexity. Snowpipe (B) is for continuous data ingestion, not report generation. Data Marketplace (E) is unlikely to perfectly match the custom

- [illegible]

myportal.utt.edu.tt, Disposable vapes

P.S. Free & New DAA-C01 dumps are available on Google Drive shared by Prep4sures: <https://drive.google.com/open?id=1GPIGJW5UVk07QxjSea7oRJCPM0pHSalm>