

New DEA-C02 Practice Questions - DEA-C02 Latest Version



P.S. Free 2026 Snowflake DEA-C02 dumps are available on Google Drive shared by PremiumVCEDump:
<https://drive.google.com/open?id=1BeFv4PiLl-zeqBZjr6m2vOyOLORHtYSI>

To let the clients be familiar with the atmosphere and pace of the real DEA-C02 exam we provide the function of stimulating the exam. In such a way, our candidates will become more confident by practising on it. And our expert team updates the DEA-C02 Study Guide frequently to let the clients practice more. So the quality of our DEA-C02 practice materials is very high and we can guarantee to you that you will have few difficulties to pass the exam.

It is known to us that time is money, and all people hope that they can spend less time on the pass. We are happy to tell you that The SnowPro Advanced: Data Engineer (DEA-C02) exam questions from our company will help you save time. With meticulous care design, our study materials will help all customers pass their exam in a shortest time. If you buy the DEA-C02 Study Materials from our company, you just need to spend less than 30 hours on preparing for your exam, and then you can start to take the exam. We believe that you will pass your exam and get the related certification with DEA-C02 study dump.

>> New DEA-C02 Practice Questions <<

Free PDF Snowflake - DEA-C02 - Useful New SnowPro Advanced: Data Engineer (DEA-C02) Practice Questions

An SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) practice questions is a helpful, proven strategy to crack the Snowflake DEA-C02 exam successfully. It helps candidates to know their weaknesses and overall performance.

PremiumVCEDump has hundreds of SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) exam dumps that are useful to practice in real time. The Snowflake DEA-C02 practice questions have a close resemblance with the actual DEA-C02 exam.

Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Sample Questions (Q285-Q290):

NEW QUESTION # 285

You are tasked with implementing a data governance strategy in Snowflake for a large data warehouse. Your objective is to classify sensitive data columns, such as customer phone numbers and email addresses, using tags. You want to define a flexible tagging system that allows different levels of sensitivity (e.g., 'Confidential', 'Restricted') to be applied to various columns. Furthermore, you need to ensure that any data replicated to different regions maintains these classifications. Which of the following statements accurately describe best practices for implementing and maintaining data classification using tags in Snowflake, especially in a multi-region setup? Choose TWO.

- A. Create a scheduled task that automatically identifies sensitive data based on regular expressions and applies the appropriate tags. This automates the classification process.

- B. Define tag schemas at the account level and replicate them to all regions. This ensures consistency of tag definitions across the entire organization.
- C. Tags and tag values must be uniquely defined across all schemas to avoid conflicts and ensure accurate data classification; Snowflake enforces uniqueness implicitly.
- D. When replicating data between regions, the tags are automatically replicated along with the data, provided that replication is configured using database replication or failover groups including the tagging schema.
- E. Always grant the ACCOUNTADMIN role to users who need to apply tags. This simplifies the process and ensures they have all necessary privileges.

Answer: B,D

Explanation:

Defining tag schemas at the account level (Option A) ensures consistency in tag definitions across the entire Snowflake account, including all regions. This is a best practice for managing tags in a multi-region environment. When replicating data between regions (Option C) using database replication or failover groups, the tags are automatically replicated along with the data, assuming the tagging schema is included in the replication configuration. Option B describes a valid approach to tag application automation, but it isn't a core best practice related to multi-region replication and tag management. Option D is incorrect because granting the ACCOUNTADMIN role provides excessive privileges and is not a recommended practice. Option E is incorrect because tag names need only be unique within their schema.

NEW QUESTION # 286

You have a Snowflake view that joins three large tables: ORDERS, CUSTOMERS, and PRODUCTS. The query accessing this view is frequently used but performs poorly. You suspect inefficient join processing and potential skew in the data¹. Which of the following strategies can be used to optimize the view's performance? (Select all that apply)

- A. Replace the view with a materialized view to precompute and store the results.
- B. Partition the underlying tables based on the join keys to improve data locality.
- C. Increase the virtual warehouse size to provide more resources for query processing. Convert the view into a table using CREATE TABLE AS SELECT (CTAS).
- D. Use JOIN hints, such as BROADCAST or MERGE, to guide the query optimizer on join strategies.
- E. Analyze the query profile to identify bottlenecks and potential data skew issues, and then re-cluster the underlying tables based on the most frequently used join keys.

Answer: A,D,E

Explanation:

Materialized views (A) can significantly improve performance by precomputing the results. JOIN hints (B) help the query optimizer choose the most efficient join strategy. Analyzing the query profile and re-clustering (C) addresses potential data skew and inefficient join processing. Increasing warehouse size (D) can help, but it's not a targeted solution for join performance. Partitioning (E) isn't directly supported in Snowflake; clustering is the analogous approach.

NEW QUESTION # 287

A data engineering team has created a Snowflake Listing to share their company's sales data¹. They want to allow consumers to access the Listing programmatically. The consumers need to know when new versions of the Listing are available. What is the MOST efficient method for consumers to be notified about new Listing versions without continuously polling Snowflake?

- A. The data consumer implements a custom webhook endpoint that the data provider calls whenever a new Listing version is published. The data provider will need to maintain a list of all consumers.
- B. The consumer sets up a scheduled task in Snowflake to periodically query the 'SHOW VERSIONS' command for the Listing.
- C. Snowflake automatically notifies consumers via an event notification service (e.g., AWS SNS, Azure Event Grid) integrated with the Data Marketplace when a new version is available.
- D. The data provider manually sends email notifications to each consumer whenever a new version is created.

Answer: C

Explanation:

Snowflake's integration with event notification services like AWS SNS or Azure Event Grid provides the most efficient and scalable way to notify consumers about new Listing versions. This eliminates the need for manual intervention or continuous polling, making it

a push-based system. Snowflake automatically handles the notification process based on preconfigured event triggers.

NEW QUESTION # 288

You are designing a data pipeline in Snowflake that involves several tasks chained together. One of the tasks, 'task_B', depends on the successful completion of 'task_A'. 'task_B' occasionally fails due to transient network issues. To ensure the pipeline's robustness, you need to implement a retry mechanism for 'task_B' without using external orchestration tools. What is the MOST efficient way to achieve this using native Snowflake features, while also limiting the number of retries to prevent infinite loops and excessive resource consumption? Assume the task definition for 'task_B' is as follows:

- A. Embed the retry logic directly within the stored procedure called by 'task_B'. The stored procedure should catch exceptions related to network issues, introduce a delay using 'SYSTEM\$WAIT', and retry the main logic. Implement a loop with a maximum retry count.
- B. Modify the task definition of 'task_B' to include a SQL statement that checks for the success of 'task_R' in the 'TASK_HISTORY' view before executing the main logic. If 'task_A' failed, use 'SYSTEM\$WAIT' to introduce a delay and then retry the main logic. Implement a counter to limit the number of retries.
- C. Leverage Snowflake's event tables like 'QUERY_HISTORY' and 'TASK_HISTORY' in the 'ACCOUNT_USAGE' schema joined with custom metadata tags to correlate specific transformation steps to execution times and resource usage. Also set up alerting based on defined performance thresholds.
- D. Utilize Snowflake's external functions to call a retry service implemented in a cloud function (e.g., AWS Lambda or Azure Function). The external function will handle the retry logic and update the task status in Snowflake.
- E. Create a separate task, 'task_C', that is scheduled to run immediately after 'task_B' will check the status of 'task_BS' in the 'TASK_HISTORY' view. If 'task_B' failed, 'task_C' will re-enable 'task_B' and suspend itself. Use the parameter on 'task_B' to limit the number of retries.

Answer: A

Explanation:

Option C is the most efficient and self-contained approach using native Snowflake features. Embedding the retry logic within the stored procedure called by 'task_B' allows for fine-grained control over the retry process, exception handling, and delay implementation. The retry count limit prevents infinite loops. Option A, while technically feasible, involves querying the 'TASK_HISTORY' view, which can be less efficient. Option B requires creating and managing an additional task. Option D introduces external dependencies, making the solution more complex. Option E does not address the retry mechanism.

NEW QUESTION # 289

You have a 'SALES' table and a 'PRODUCTS' table. The 'SALES' table contains daily sales transactions, including 'SALE_DATE', 'PRODUCT_ID', and 'QUANTITY'. The 'PRODUCTS' table contains 'PRODUCT' and 'CATEGORY'. You need to create a materialized view to track the total quantity sold per category daily, optimized for fast query performance. You anticipate frequent updates to the 'SALES' table but infrequent changes to the 'PRODUCTS' table. Which of the following strategies would provide the MOST efficient materialized view implementation, considering both data freshness and query performance?

- A. Create a standard materialized view that joins 'SALES' and 'PRODUCTS', grouping by 'SALE_DATE' and 'CATEGORY', and defining a clustering key on 'CATEGORY'.
- B. Create a standard materialized view that joins 'SALES' and 'PRODUCTS', grouping by 'SALE_DATE' and 'CATEGORY', and defining a clustering key on 'SALE_DATE' and 'CATEGORY'.
- C. Create a standard materialized view that joins 'SALES' and 'PRODUCTS', grouping by 'SALE_DATE' and 'CATEGORY' without any specific clustering key.
- D. Create two materialized views: one for daily sales by product and another joining the first with 'PRODUCTS' to aggregate by category. Cluster the first view by 'SALE_DATE' and the second by 'CATEGORY'.
- E. Create a standard materialized view that joins 'SALES' and 'PRODUCTS', grouping by 'SALE_DATE' and 'CATEGORY', and defining a clustering key on 'SALE_DATE'.

Answer: E

Explanation:

Option B is most efficient. Clustering the materialized view on 'SALE_DATE' will significantly improve query performance when filtering or grouping by date, which is a common operation in time-series data. Although frequent updates will affect the maintenance costs of the materialized view, querying on date will be very efficient. Option A is less efficient due to the lack of clustering. Option C may not be the best choice if filtering/grouping primarily occurs on date. Option D is also good, but Option B is better if most of the query filter is on SALE_DATE. Option E introduces complexity and two refreshes may create a delay in data available.

NEW QUESTION # 290

High as 98 to 100 percent of exam candidates pass the exam after refer to the help of our DEA-C02 practice braindumps. So DEA-C02 study guide is high-effective, high accurate to succeed. That is the reason why we make it without many sales tactics to promote our DEA-C02 Learning Materials, their brand is good enough to stand out in the market. Download our DEA-C02 training prep as soon as possible and you can begin your review quickly.

DEA-C02 Latest Version: <https://www.premiumvcedump.com/Snowflake/valid-DEA-C02-premium-vce-exam-dumps.html>

Please stop, and pay attention to our DEA-C02 prep training, Secure payment, Snowflake New DEA-C02 Practice Questions
Dear friends, as you know, there are some fateful exams which can decide your destiny when you compete with others on working environment or job fair, because certificates are indispensable gauges for boss to estimate your capacity, Snowflake New DEA-C02 Practice Questions Therefore, you have no need to worry about the types of your cellphone.

His quest for doing everything digitally has defined Centric Digital's operations New DEA-C02 Practice Questions and fully removed the use of paper from his work or personal life, Lizard straightens you out by showing you how to set up an automated repository.

Useful Snowflake New DEA-C02 Practice Questions Are Leading Materials & First-Grade DEA-C02 Latest Version

Please stop, and pay attention to our DEA-C02 prep training. Secure payment, Dear friends, as you know, there are some fateful exams which can decide your destiny when you compete with others on working environment DEA-C02 or job fair, because certificates are indispensable gauges for boss to estimate your capacity.

Therefore, you have no need to worry about DEA-C02 Latest Version the types of your cellphone, You can contact us whenever you need help.

DOWNLOAD the newest PremiumVCEDump DEA-C02 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1BeFv4PiLl-zeqBZjr6m2vOyOLORHfYS!>