

Training DEA-C02 Pdf Exam Pass at Your First Attempt | Snowflake DEA-C02: SnowPro Advanced: Data Engineer (DEA-C02)



Before we start develop a new DEA-C02 study materials, we will prepare a lot of materials. After all, we must ensure that all the questions and answers of the DEA-C02 study materials are completely correct. First of all, we have collected all relevant reference books. Most of the DEA-C02 Study Materials are written by the famous experts in the field. They are widely read and accepted by people. Through careful adaption and reorganization, all knowledge will be integrated in our DEA-C02 study materials.

So rest assured that with the SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) practice questions you will not only make the entire DEA-C02 exam dumps preparation process and enable you to perform well in the final SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) certification exam with good scores. To provide you with the updated Snowflake DEA-C02 Exam Questions the Snowflake offers three months updated SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) exam dumps download facility, Now you can download our updated DEA-C02 practice questions up to three months from the date of SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) exam purchase.

[**>> Training DEA-C02 Pdf <<**](#)

Training DEA-C02 Pdf - 100% Pass 2026 First-grade DEA-C02: SnowPro Advanced: Data Engineer (DEA-C02) Latest Dumps Questions

The profession of our experts is expressed in our DEA-C02 training prep thoroughly. They are great help to catch on the real knowledge of DEA-C02 exam and give you an unforgettable experience. Do no miss this little benefit we offer for we give some discounts on our DEA-C02 Exam Questions from time to time though the price of our DEA-C02 study guide is already favourable. And every detail of our DEA-C02 learning braindumps is perfect!

Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Sample Questions (Q138-Q143):

NEW QUESTION # 138

A data engineering team is building a data pipeline in Snowflake. They are using tasks and streams to incrementally load data into a fact table. The team needs to monitor the pipeline's performance and ensure data lineage. What are the valid and most effective techniques to ensure that this pipeline adheres to compliance and governance rules?

- A. Enable Snowflake Horizon features, which include Data Lineage, Object Dependencies and Discovery and integrate it with the data lake and also tag the data pipeline.
- B. Use Account Usage views like 'TASK HISTORY' and 'STREAM_LAG' to track task execution and stream latency, create stored procedures to log metadata about each pipeline run to a separate metadata table, and rely on developers to manually document the pipeline's data flow and policy enforcement.
- C. Leverage Snowflake's replication features for disaster recovery, monitor only the replication lag, and disable all security policies to improve performance since those tasks have already been validated during the initial deployment of the software.

- D. Use a third-party data catalog to track lineage, monitor task performance via 'TASK_HISTORY', and ignore data masking and row-level security policies for simplicity in the initial implementation.
- E. Implement Snowflake's Data Lineage and Object Dependencies features to track data flow automatically, create Alerts based on 'TASK HISTORY' to monitor task failures, and enforce data masking and row-level security policies at the table level. Use Snowflake's tags to categorise and classify objects.

Answer: A,E

Explanation:

Option B and E offers the most comprehensive approach. Snowflake's Data Lineage and Object Dependencies features, combined with Alerts based on 'TASK HISTORY', offer automated monitoring and data flow tracking. Enforcing data masking and row-level security is crucial for data governance and compliance, and tagging enables easy categorisation and discovery. Though Horizon and Data Lineage offers data flow tracking, alerting for task failures, data masking and row-level security policies and tagging can be integrated together. Options A lacks automated lineage tracking and relies on manual documentation, which is error-prone. Option C ignores crucial security policies, which is unacceptable. Option D focuses only on disaster recovery and neglects security and monitoring aspects.

NEW QUESTION # 139

You have a table 'EVENTS' containing application event data with columns 'EVENT ID', 'USER ID', 'EVENT TYPE', and EVENT DETAILS (VARCHAR). The 'EVENT DETAILS' column contains comma-separated key-value pairs (e.g., 'location=USA,device=mobile,os=iOS'). Your objective is to transform this structured data into a VARIANT column named 'EVENT JSON' in a new table 'EVENTS JSON'. The data in EVENT DETAILS has inconsistent key-value pairs across different rows. Which of the following methods are the most efficient and scalable to parse the key-value pairs in 'EVENT DETAILS' and construct the JSON objects?

- A. Use only REGEXP EXTRACT ALL' with appropriate regular expressions to extract all keys and values into arrays, then use a JavaScript UDF to combine them into a JSON object.
- B. Use a Java UDF that iterates through the string, splitting it based on commas and equals signs, and then constructs a JSON object using a JSON library.
- C. Use 'SPLIT to split the key-value pairs into an array, then use a LATERAL FLATTEN to create rows from array, then use 'SPLIT again to split each row by '='. Finally, construct the JSON using 'OBJECT CONSTRUCT'.
- D. Utilize to split the key-value pairs into rows, then use 'REGEXP_EXTRACT to extract the key and value. Finally, use 'OBJECT_CONSTRUCT and to construct the JSON object.
- E. Use a combination of 'SPLIT', 'REGEXP_REPLACE and 'OBJECT_CONSTRUCT within a user-defined function (UDF) to parse the string and build the JSON object.

Answer: C,D

Explanation:

Options B and E are the most efficient because they use Snowflake's built-in functions for parsing the data before constructing the JSON object. 'SPLIT TO_TABLE' (B) and LATERAL FLATTEN (E) are the more scalable way, if the number of key-value pairs might be high in string. While UDFs (A, C, and D) can achieve the transformation, they can be less performant and harder to maintain compared to SQL-based solutions. 'REGEXP EXTRACT ALL' is less efficient in this case as you would still need to iterate through the arrays it creates. Splitting data with 'SPLIT' and utilizing LATERAL FLATTEN provides a good balance between performance and readability for handling semi-structured data within Snowflake's SQL environment.

NEW QUESTION # 140

You are designing a data sharing solution for a multi-tenant application where each tenant's data must be isolated. You have a 'sales' table with a 'tenant_id' column. You need to implement row-level security to ensure that each tenant can only access their own data when querying the shared table. Which of the following approaches, considering performance and security, is the MOST suitable for implementing this row-level filtering in Snowflake?

- A. Implement a row access policy on the 'sales' table that filters data based on the 'tenant_id' column and the current role or user context.
- B. Create a separate VIEW for each tenant, filtering by 'tenant_id'. Grant each tenant access only to their respective view.
- C. Implement a user-defined function (UDF) that checks the current user's tenant ID and returns a boolean value indicating whether the row should be visible. Use this UDF in a WHERE clause in every query.
- D. Use Snowflake's data masking policies to mask all data for tenants other than the one currently querying the table.
- E. Create a scheduled task that duplicates the sales table into a new table for each tenant, filtering by the tenant_id.

Answer: A

Explanation:

Row access policies are the most efficient and scalable solution for row-level security in Snowflake. They are applied at the table level and automatically enforced for all queries, ensuring consistent data isolation. Creating separate views for each tenant is administratively burdensome and doesn't scale well. UDFs can impact performance. Data masking policies do not filter rows, they only redact data. Creating scheduled task create data duplication, it's not ideal for data movements.

NEW QUESTION # 141

You are designing a data loading process for a high-volume streaming data source. The data arrives as Avro files in an AWS S3 bucket. You need to load this data into a Snowflake table with minimal latency and operational overhead. Which of the following combinations of Snowflake features and configurations would be MOST suitable for this scenario? (Select TWO)

- A. Configure an external table pointing to the S3 bucket and query the Avro files directly from Snowflake.
- B. Use a Kafka connector to stream data directly from the Kafka topic to Snowflake.
- C. Use the 'COPY INTO' command with a scheduled task that runs every 5 minutes to load new files from the S3 bucket.
- D. Create a custom Spark application that reads Avro files from S3, transforms the data, and then writes it to Snowflake using the Snowflake Spark connector.
- E. Implement Snowpipe with auto-ingest configured to listen for S3 event notifications whenever a new Avro file is added to the bucket.

Answer: B,E

Explanation:

Explanation: Options B and C offer the best combination of low latency and operational efficiency for streaming data. Snowpipe with auto-ingest provides near real-time loading triggered by S3 events. A Kafka connector provides a direct data stream to Snowflake. Option A introduces latency due to the scheduled task interval and doesn't scale well for high-volume streams. Option D adds operational overhead with Spark application management. Option E is suitable for ad-hoc querying but not ideal for continuous data loading.

NEW QUESTION # 142

You are tasked with optimizing a data pipeline that loads data from an external cloud storage location into Snowflake, transforms it, and then loads it into reporting tables. The pipeline is experiencing intermittent performance issues. You want to proactively identify and address these issues. Which of the following monitoring techniques and Snowflake features would be MOST effective for continuous monitoring and performance optimization?

- A. Rely solely on Snowflake's default query history and resource monitors. These automatically track performance and usage, providing sufficient insight without additional configuration.
- B. Utilize Snowflake's System Functions to periodically query performance views (e.g., 'QUERY_HISTORY,' and write aggregated metrics to a dedicated monitoring table. Configure a scheduled task to generate alerts based on predefined thresholds.
- C. Implement custom logging and monitoring using Snowflake Scripting and User-Defined Functions (UDFs) to capture granular performance metrics at each stage of the pipeline and push notifications via external functions to a monitoring service.
- D. Focus exclusively on optimizing SQL queries and data transformations. Monitoring is unnecessary since Snowflake automatically handles performance optimization.
- E. Enable Snowflake's Auto-Suspend and Auto-Resume features on the warehouse. This is the most efficient way to manage resources and optimize costs, indirectly addressing performance concerns.

Answer: B,C

Explanation:

Options B and C provide the most effective methods for continuous monitoring and performance optimization. Option B allows for highly customized and granular monitoring of the entire pipeline, enabling proactive issue identification through external notifications. Option C leverages Snowflake's built-in system functions and task scheduling to create a robust monitoring and alerting system. Option A is insufficient as default monitoring may not provide the granularity needed. Option D is incorrect because monitoring is crucial. Option E primarily focuses on cost optimization, not performance monitoring.

NEW QUESTION # 143

The Internet is increasingly becoming a platform for us to work and learn, while many products are unreasonable in web design, and too much information is not properly classified. It's disorganized. Our DEA-C02 study materials draw lessons from the experience of failure, will all kinds of qualification examination has carried on the classification of clear layout, at the same time the user when they entered the DEA-C02 Study Materials page in the test module classification of clear, convenient to use a very short time to find what they want to study, which began the next exercise.

DEA-C02 Latest Dumps Questions: <https://www.actaltorrent.com/DEA-C02-questions-answers.html>

Here DEA-C02 study dumps will drag you from the confusion, With the number of people who take the exam increasing, the DEA-C02 exam has become more and more difficult for many people, Snowflake Training DEA-C02 Pdf Teachers and educationist have tried many ways to solve this problem, Our DEA-C02 Exam Preparation Material provides you everything you will need to take a DEA-C02 certifications examination, The practice software of the Snowflake DEA-C02 exam will help you to get ready for the DEA-C02 exam in the short time and you will feel confident while taking the actual SnowPro Advanced: Data Engineer (DEA-C02) Exam DEA-C02 exam

Processing default locales, Make the time you spend taking photos together fun, Here DEA-C02 Study Dumps will drag you from the confusion, With the number of people who take the exam increasing, the DEA-C02 exam has become more and more difficult for many people.

Snowflake DEA-C02 dumps VCE file - Testking DEA-C02 real dumps

Teachers and educationist have tried many ways to solve this problem, Our DEA-C02 Exam Preparation Material provides you everything you will need to take a DEA-C02 certifications examination.

The practice software of the Snowflake DEA-C02 exam will help you to get ready for the DEA-C02 exam in the short time and you will feel confident while taking the actual SnowPro Advanced: Data Engineer (DEA-C02) Exam DEA-C02 exam.

