

Pass Guaranteed Snowflake DEA-C02 - SnowPro Advanced: Data Engineer (DEA-C02) Marvelous Exam Papers



What's more, part of that Getcertkey DEA-C02 dumps now are free: <https://drive.google.com/open?id=10c8QSCqRCNPIPSmVV52jkkBhTpNqIJ>

For candidates who will attend the exam, some practice is necessary. DEA-C02 exam materials are valid and high-quality. We have a professional team to search for the first-hand information for the exam. We also have strict requirements for the questions and answers of DEA-C02 exam materials, we ensure you that the DEA-C02 Training Materials are most useful tool, which can help you pass the exam just one time. In addition, we offer you free update for one year after purchasing, we also have online service stuff, if you have any questions, just contact us.

Are you ready to accept this challenge? Looking for the simple, quick, and easiest way to pass the career advancement SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) certification exam? If your answer is yes then you do not need to worry about it. Just visit the Getcertkey and explore the top features of SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) exam practice test questions offered by the trusted platform Getcertkey. With Getcertkey DEA-C02 Dumps questions you can easily prepare well and feel confident to pass the final SnowPro Advanced: Data Engineer (DEA-C02) exam easily.

>> DEA-C02 Exam Papers <<

Snowflake DEA-C02 Pass4sure Study Materials, DEA-C02 Real Questions

All of these prep formats pack numerous benefits necessary for optimal preparation. This SnowPro Advanced: Data Engineer (DEA-C02) (DEA-C02) practice material contains actual Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Questions that invoke conceptual thinking. Getcertkey provides you with free-of-cost demo versions of the product so that you may check the validity and actuality of the Snowflake DEA-C02 Dumps PDF before even buying it. We also offer a money-back guarantee, which means we are obliged to return 100% of your sum (terms and conditions apply) in case of any unsatisfactory results.

Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Sample Questions (Q185-Q190):

NEW QUESTION # 185

A large e-commerce company is experiencing performance issues with its daily sales report queries. These queries aggregate data from a fact table 'SALES FACT' (100 billion rows) and several dimension tables, including 'CUSTOMER DIM', 'PRODUCT DIM', and 'DATE DIM'. The queries are run every morning and are essential for business decision-making. The team has identified that the 'SALES FACT' table's primary key is 'SALE ID', but the queries frequently filter and join on 'CUSTOMER' and 'PRODUCT ID'. You want to use query acceleration service for these reports without changing query logic. Which combination of actions will MOST effectively leverage query acceleration service, assuming sufficient credits?

- A. Increase the size of the virtual warehouse used for running the reports and enable query acceleration. Set the parameter to a high value.

- B. Create materialized views that pre-aggregate the sales data based on 'CUSTOMER ID', 'PRODUCT ID, and 'DATE ID, then enable query acceleration on the virtual warehouse.
- C. Enable clustering on the 'CUSTOMER DIM' and 'PRODUCT DIMS tables.
- D. Enable Automatic Clustering on the 'SALES FACT table based on 'CUSTOMER ID' and 'PRODUCT ID, then enable query acceleration on the virtual warehouse.
- E. Enable search optimization on the columns 'CUSTOMER ID' and 'PRODUCT ID of the 'SALES FACT table, then enable query acceleration on the virtual warehouse. Set the QUERY_ACCELERATION_MAX_SCALE_FACTOR parameter to a reasonable value based on testing.

Answer: E

Explanation:

Enabling search optimization on and in the 'SALES FACT table is the most effective approach for query acceleration in this scenario. This allows Snowflake to efficiently find the rows needed for the reports without scanning the entire table. Automatic Clustering improves data organization, but it doesn't directly accelerate individual queries in the same way. Materialized views are also useful, but require additional storage and maintenance. Simply increasing the warehouse size and enabling query acceleration without addressing the data organization or indexing might not be as cost-effective. Clustering dimension tables won't affect the performance on the large fact table.

NEW QUESTION # 186

You are planning to monetize a dataset on the Snowflake Marketplace. You want to provide potential customers with sample data to evaluate before they purchase a full subscription. Which of the following strategies are valid and recommended for offering a free sample of your data within the Snowflake Marketplace? (Select all that apply)

- A. Upload a sample CSV file to a publicly accessible S3 bucket and provide the link in the Marketplace listing description. Consumers can download and load this data into their own Snowflake account for evaluation.
- B. Provide the consumer with the script to create a database link to your data, allowing them read-only access to a pre-defined sample table, and then revoke the access after a set period.
- C. Create a separate share containing a subset (e.g., a smaller number of rows or columns) of the full dataset and offer this share as a free trial listing on the Marketplace.
- D. Create a view that filters the dataset based on a sampling algorithm (e.g., 'SAMPLE ROW' clause) and share the view through the Marketplace.
- E. Offer a 'free trial' subscription on the primary listing that automatically expires after a set period (e.g., 7 days), allowing customers to access the full dataset during the trial period. You will need to write custom code to manage trial expiration and data access restrictions based on the trial status.

Answer: C,D

Explanation:

Option A is valid: Creating a separate share with a subset of the data is a common and secure way to offer a free sample. Option C is valid: sharing a view filtered by a sampling algorithm offers a representative sample without requiring manual data management. Option B is less ideal because managing trial expiration and access controls requires custom coding and increases complexity. Offering consumers a database link (Option D) is a valid concept but less secure than using Snowflake's native sharing mechanisms for trial access. Option E is less integrated with the Marketplace experience; the consumer must leave Snowflake to access the sample data, and there is no tracking of who is accessing the sample.

NEW QUESTION # 187

You are working on a Snowpark Python application that needs to process a stream of data from Kafka, perform real-time aggregations, and store the results in a Snowflake table. The data stream is highly variable, with occasional spikes in traffic that overwhelm your current Snowpark setup, leading to significant latency in processing. Which of the following strategies, either individually or in combination, would be MOST effective to handle these traffic spikes and ensure near real-time processing?

- A. Configure the Snowflake warehouse used by your Snowpark application to use auto-suspend and auto-resume with a short auto-suspend time to minimize costs during periods of low traffic.
- B. Use 'CACHE RESULT for all queries in snowpark that use Kafka
- C. Implement dynamic warehouse scaling. Utilize Snowflake's Resource Monitors and the ability to programmatically resize warehouses through Snowpark. Monitor the queue depth or latency of your Snowpark application, and dynamically scale up the warehouse size when thresholds are exceeded. Then, scale it back down when traffic subsides.
- D. Implement a message queuing system (e.g., RabbitMQ, Kafka) between Kafka and your Snowpark application to buffer

incoming data during traffic spikes.

- E. Use Snowpark's async actions (e.g., to offload data processing to separate threads or processes, allowing your main Snowpark application to continue receiving data.





Answer: C,D

Explanation:

Options A and D offer the best approach. Implementing a message queue (A) provides a buffer for incoming data during spikes, preventing your Snowpark application from being overwhelmed. Dynamic warehouse scaling (D) allows you to automatically increase the compute resources available to your Snowpark application when needed, ensuring it can handle the increased workload. Auto suspend/resume (B) is good for cost optimization but doesn't address the processing capacity during spikes. Async actions (C) can help, but are not as scalable or resilient as a proper message queue combined with dynamic warehouse scaling. Caching results (E) is irrelevant since the data from Kafka is always changing.

NEW QUESTION # 188

You have a table named 'TRANSACTIONS' which is frequently queried by 'TRANSACTION_DATE' and 'CUSTOMER_ID'. You want to define a clustering strategy for this table. You are aware that defining multiple clustering keys is possible. Given the following considerations, which of the following clustering strategies would provide the BEST performance AND minimize reclustering costs, assuming both columns have similar cardinality and are equally used in WHERE clauses? (Assume cost optimization is the most critical factor if performance difference is minimal.)

- A. 
- B. Create two separate tables: one clustered by 'TRANSACTION_DATE' and another clustered by 'CUSTOMER_ID', and use appropriate views to redirect queries to the correct table.
- C. 
- D. 
- E. 

Answer: A,C

Explanation:

Clustering by both and is beneficial when queries frequently filter on both columns. Either order, TRANSACTION_DATE followed by CUSTOMER_ID, or CUSTOMER_ID followed by TRANSACTION_DATE, could provide similar performance depending on query patterns. However, having both columns as clustering keys will allow better filtering of micro-partitions. Creating two separate tables as suggested in Option E introduces complexity in data maintenance (two copies of data) and query redirection logic, increasing overall cost. While it might provide optimal performance for specific query patterns, the cost is generally higher than using a composite clustering key when both keys are frequently used.

NEW QUESTION # 189

A data engineer is using the Snowflake Spark connector to read a large table from Snowflake into a Spark DataFrame. The table contains a 'TIMESTAMP_NTZ' column. After loading the data, the engineer observes that the values in the 'TIMESTAMP_NTZ' column are not preserved accurately when retrieved from the DataFrame. What are the potential issues and what configurations can be adjusted in Snowflake to improve the result?

- ☐ The Snowflake Spark connector automatically converts 'TIMESTAMP_NTZ' to UTC. There is no configuration option to preserve the original timezone information in Spark.
- ☐ The issue is caused by a mismatch in the Spark session's timezone and the Snowflake session's timezone. Set the 'spark.sql.session.timezone' configuration in Spark to match the 'TIMEZONE' parameter set in Snowflake.
- ☐ Spark does not natively support 'TIMESTAMP_NTZ'. Convert the 'TIMESTAMP_NTZ' column to a string representation in Snowflake using 'TO_VARCHAR' before reading it into Spark. Then convert string to Timestamp object, and update session timezone if required.
- ☐ The default behavior of the Snowflake Spark connector is to convert 'TIMESTAMP_NTZ' to 'java.sql.Timestamp' which can lead to data loss due to differing precisions. To prevent this, set the 'sfTimestampNTZAsString' option to 'true' when reading data. This forces the connector to represent 'TIMESTAMP_NTZ' values as strings in the DataFrame. Then cast String to Timestamp object in a specific timezone. For example: `df.withColumn("timestamp_col", to_timestamp(col("timestamp_col"), "yyyy-MM-dd HH:mm:ss.SSS"))`
- ☐ The 'sfTimestampNTZAsString' option in the Spark Connector configuration should be used to preserve time zone information. Configure the 'SnowflakeConf' to set the parameter 'sfTimestampNTZAsString = false', which will load the timestamps without time zone conversion.

- A. Option C
- B. Option B
- C. Option D
- D. Option A
- E. Option E

Answer: B,C

Explanation:

Explanation: Option B identifies a key cause: timezone mismatches between Spark and Snowflake. Ensuring these are aligned is essential. Option D proposes converting 'TIMESTAMP NTZ to strings using 'sfTimestampNTZAsString' is the correct approach to preserve precise values, and setting the Spark session timezone is also critical. Option A is incorrect since there are parameters. Option C involves unnecessary data transformation within Snowflake which is what we are trying to avoid with the connector. Setting 'sfTimestampNTZAsString' to false causes conversion which lead to data issue.

NEW QUESTION # 190

.....

While most people would think passing Snowflake certification DEA-C02 exam is difficult. However, if you choose Getcertkey, you will find gaining Snowflake certification DEA-C02 exam certificate is not so difficult. Getcertkey training tool is very comprehensive and includes online services and after-sales service. Professional research data is our online service and it contains simulation training examination and practice questions and answers about Snowflake Certification DEA-C02 Exam. Getcertkey's after-sales service is not only to provide the latest exam practice questions and answers and dynamic news about Snowflake DEA-C02 certification, but also constantly updated exam practice questions and answers and binding.

DEA-C02 Pass4sure Study Materials: https://www.getcertkey.com/DEA-C02_braindumps.html

If you are determined to learn something, our DEA-C02 test torrent material will be your best choice, Choose the nay type of DEA-C02 practice exam questions that fit your DEA-C02 exam preparation requirement and budget and start preparation without wasting further time, Our DEA-C02 learning questions engage our working staff in understanding customers' diverse and evolving expectations and incorporate that understanding into our strategies, thus you can 100% trust our DEA-C02 exam engine, If you want to pass exam and get the related certification in the shortest time, the DEA-C02 DEA-C02 study materials from our company will be your best choice.

Implement Measures and Measure Groups in a Cube, Person Responsible for Quality, If you are determined to learn something, our DEA-C02 Test Torrent material will be your best choice.

Choose the nay type of DEA-C02 practice exam questions that fit your DEA-C02 exam preparation requirement and budget and start preparation without wasting further time.

2026 DEA-C02 – 100% Free Exam Papers | Excellent DEA-C02 Pass4sure Study Materials

Our DEA-C02 learning questions engage our working staff in understanding customers' diverse and evolving expectations and incorporate that understanding into our strategies, thus you can 100% trust our DEA-C02 exam engine.

If you want to pass exam and get the related certification in the shortest time, the DEA-C02 DEA-C02 study materials from our company will be your best choice.

How can I get discounts?

- Snowflake Offers Valid and Real Snowflake DEA-C02 Exam Questions □ Open “ www.examcollectionpass.com ” and search for 「 DEA-C02 」 to download exam materials for free □ DEA-C02 Exam Quiz
- Knowledge DEA-C02 Points □ Popular DEA-C02 Exams □ Test DEA-C02 Dumps □ Search for □ DEA-C02 □ and obtain a free download on ► www.pdfvce.com ◀ □ Updated DEA-C02 Test Cram
- Snowflake Offers Valid and Real Snowflake DEA-C02 Exam Questions □ Download 《 DEA-C02 》 for free by simply searching on ► www.validtorrent.com □ □ Guaranteed DEA-C02 Questions Answers
- 100% Pass Quiz Snowflake - DEA-C02 - Authoritative SnowPro Advanced: Data Engineer (DEA-C02) Exam Papers □ Open □ www.pdfvce.com □ enter 「 DEA-C02 」 and obtain a free download □ Popular DEA-C02 Exams
- Reliable DEA-C02 Test Prep □ DEA-C02 Reliable Test Book □ Free DEA-C02 Vce Dumps □ Simply search for (DEA-C02) for free download on ✓ www.vce4dumps.com □ ✓ □ □ Latest DEA-C02 Test Sample
- 2026 Realistic DEA-C02 Pass4sure Study Materials - Snowflake SnowPro Advanced: Data Engineer (DEA-C02) Exam Papers 100% Pass □ Simply search for 【 DEA-C02 】 for free download on ► www.pdfvce.com □ □ Detail DEA-C02 Explanation
- DEA-C02 Exam Papers 100% Pass | Latest DEA-C02 Pass4sure Study Materials: SnowPro Advanced: Data Engineer (DEA-C02) □ Easily obtain □ DEA-C02 □ for free download through { www.prepawaypdf.com } □ DEA-C02 Reliable Test Book
- Latest DEA-C02 Test Sample □ Latest DEA-C02 Exam Question □ DEA-C02 Valid Braindumps Free □ Search for

