New Snowflake ARA-C01 Dumps Ebook - ARA-C01 Test Questions Answers

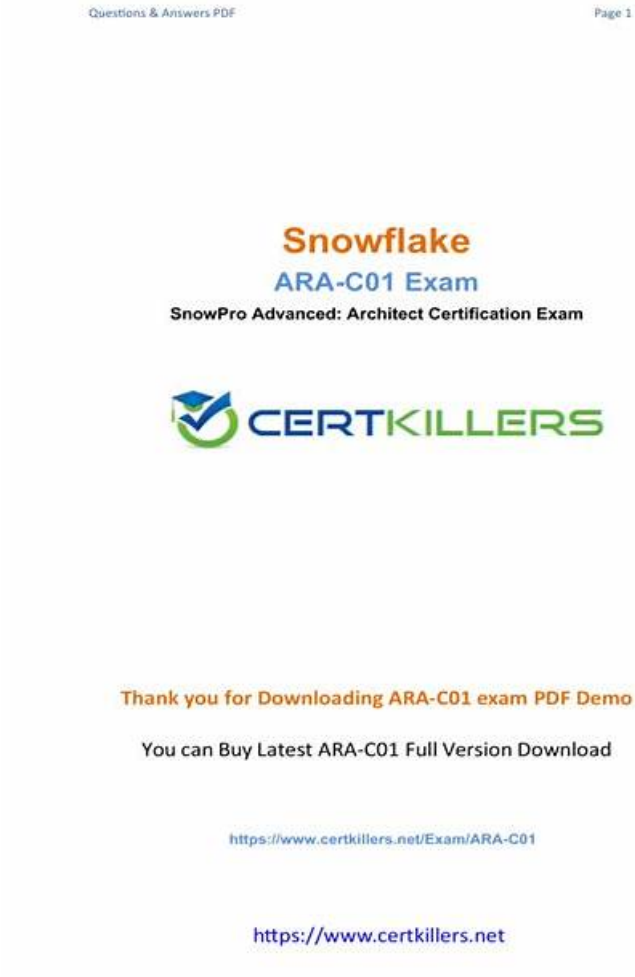

What's more, part of that PracticeDump ARA-C01 dumps now are free: https://drive.google.com/open?id=1-iCn-ESOpcWlOKzzfvCZnj4eDvC95vzT

I am glad to introduce a secret weapon for all of the candidates to pass the exam as well as get the related certification without any more ado-- our ARA-C01 study materials. You can only get the most useful and efficient study materials with the most affordable price. With our ARA-C01 practice test, you only need to spend 20 to 30 hours in preparation since there are all essence contents in our ARA-C01 Study Materials. What's more, if you need any after service help on our ARA-C01 exam guide, our after service staffs will always offer the most thoughtful service for you.

Snowflake ARA-C01 Certification Exam is not for the faint-hearted. It is a rigorous and challenging exam that requires a deep understanding of Snowflake architecture, data modeling, performance optimization, security, and administration. ARA-C01 exam consists of 60 multiple-choice questions that must be completed within 120 minutes. The passing score for the ARA-C01 exam is 80%, and candidates who pass the exam are awarded the SnowPro Advanced Architect Certification.

Snowflake ARA-C01: SnowPro Advanced Architect Certification Exam is an essential step for professionals looking to advance their careers in the field of data warehousing and cloud computing. It is designed to test an individual's knowledge and expertise in advanced Snowflake architecture and is a valuable accomplishment that demonstrates an individual's commitment to the platform.

>> New Snowflake ARA-C01 Dumps Ebook <<

Free PDF Snowflake - ARA-C01 High Hit-Rate New Dumps Ebook

In order to show you how efficient our ARA-C01 exam dump is, we allow you to download a demo version for free! You will have a chance to peak into the program and then make your final purchase decision. We are absolutely sure that once you see what's inside, you will buy it immediately without any hesitation! ARA-C01 Exam Dump also provide customer service, in case you have any inquiry or question, our professional Customer Support will be available for you 24/7. 365 days a Year.

Snowflake SnowPro Advanced Architect Certification Sample Questions (Q159-Q164):

NEW QUESTION # 159
What is a key consideration when setting up search optimization service for a table?

- A. The table must be clustered with a key having multiple columns for effective search optimization.
- B. Search optimization service can help to optimize storage usage by compressing the data into a GZIP format.
- C. Search optimization service can significantly improve query performance on partitioned external tables.
- D. Search optimization service works best with a column that has a minimum of 100 K distinct values.

Answer: D

Explanation:
Search optimization service is a feature of Snowflake that can significantly improve the performance of certain types of lookup and analytical queries on tables. Search optimization service creates and maintains a persistent data structure called a search access path, which keeps track of which values of the table's columns might be found in each of its micro-partitions, allowing some micro-partitions to be skipped when scanning the table1.
Search optimization service can significantly improve query performance on partitioned external tables, which are tables that store data in external locations such as Amazon S3 or Google Cloud Storage. Partitioned external tables can leverage the search access path to prune the partitions that do not contain the relevant data, reducing the amount of data that needs to be scanned and transferred from the external location2.
The other options are not correct because:
* A. Search optimization service works best with a column that has a high cardinality, which means that the column has many distinct values. However, there is no specific minimum number of distinct values required for search optimization service to work effectively. The actual performance improvement depends on the selectivity of the queries and the distribution of the data1.
* C. Search optimization service does not help to optimize storage usage by compressing the data into a GZIP format. Search optimization service does not affect the storage format or compression of the data, which is determined by the file format options of the table. Search optimization service only creates an additional data structure that is stored separately from the table data1.
* D. The table does not need to be clustered with a key having multiple columns for effective search optimization. Clustering is a feature of Snowflake that allows ordering the data in a table or a partitioned external table based on one or more clustering keys. Clustering can improve the performance of queries that filter on the clustering keys, as it reduces the number of micro-partitions that need to be scanned. However, clustering is not required for search optimization service to work, as search optimization service can skip micro-partitions based on any column that has a search access path, regardless of the clustering key3.
References:
* 1: Search Optimization Service | Snowflake Documentation
* 2: Partitioned External Tables | Snowflake Documentation
* 3: Clustering Keys | Snowflake Documentation

NEW QUESTION # 160
A global company needs to securely share its sales and Inventory data with a vendor using a Snowflake account.
The company has its Snowflake account In the AWS eu-west 2 Europe (London) region. The vendor's Snowflake account Is on the Azure platform in the West Europe region. How should the company's Architect configure the data share?

- A. 1. Create a share.2. Create a reader account for the vendor to use.3. Add the reader account to the share.
- B. 1. Create a new role called db_share.2. Grant the db_share role privileges to read data from the company database and schema.3. Create a user for the vendor.4. Grant the ds_share role to the vendor's users.
- C. 1. Promote an existing database in the company's local account to primary.2. Replicate the database to Snowflake on Azure in the West-Europe region.3. Create a share and add objects to the share.4. Add a consumer account to the share for the vendor to access.
- D. 1. Create a share.2. Add objects to the share.3. Add a consumer account to the share for the vendor to access.

Answer: D

Explanation:

The correct way to securely share data with a vendor using a Snowflake account on a different cloud platform and region is to create a share, add objects to the share, and add a consumer account to the share for the vendor to access. This way, the company can control what data is shared, who can access it, and how long the share is valid. The vendor can then query the shared data without copying or moving it to their own account.
The other options are either incorrect or inefficient, as they involve creating unnecessary reader accounts, users, roles, or database replication.
https://learn.snowflake.com/en/certifications/snowpro-advanced-architect/

NEW QUESTION # 161
An Architect runs the following SQL query:
How can this query be interpreted?

- A. FILEROWS is a file. FILE_ROW_NUMBER is the file format location.
- B. FILEROWS is a stage. FILE_ROW_NUMBER is line number in file.
- C. FILEROWS is the table. FILE_ROW_NUMBER is the line number in the table.
- D. FILERONS is the file format location. FILE_ROW_NUMBER is a stage.

Answer: B

NEW QUESTION # 162
Files arrive in an external stage every 10 seconds from a proprietary system. The files range in size from 500 K to 3 MB. The data must be accessible by dashboards as soon as it arrives.
How can a Snowflake Architect meet this requirement with the LEAST amount of coding? (Choose two.)

- A. Use a materialized view on an external table.
- B. Use a combination of a task and a stream.
- C. Use Snowpipe with auto-ingest.
- D. Use the COPY INTO command.
- E. Use a COPY command with a task.

Answer: B,C

NEW QUESTION # 163
An Architect is defining transaction rules to adhere to ACID properties to ensure that executed statements are either committed or rolled back. Based on this scenario, what characteristics of transactions should be considered? (Select TWO).

- A. The autocommit setting can be changed inside a stored procedure.
- B. An explicit transaction can be started by executing a BEGIN TRANSACTION statement and can be ended by executing an END TRANSACTION statement.
- C. Explicit transactions should contain DDL, DML, and query statements.
- D. Explicit transactions should contain only DML statements and query statements. All DDL statements implicitly commit active transactions.
- E. An explicit transaction can be started by executing a BEGIN WORK statement and can be ended by executing a COMMIT WORK statement.

Answer: D,E

Explanation:
Comprehensive and Detailed 150 to 250 words of Explanation From Snowflake SnowPro Architect exam scope and all publicly documented material:
Snowflake supports transactional behavior for DML (and SELECT statements within a transaction context), but DDL statements have special behavior: many DDL operations implicitly commit the current transaction.
Because of that, mixing DDL with DML inside an explicit transaction undermines the "all-or-nothing" rollback expectations and complicates ACID-driven rules. Therefore, a key design characteristic is that explicit transactions should generally be limited to DML and query statements, while recognizing that DDL can implicitly commit and break transactional boundaries (Choice C).
Additionally, Snowflake supports explicit transaction control statements using standard SQL forms; BEGIN WORK and COMMIT WORK are valid ways to open and close an explicit transaction (Choice D). From an architecting standpoint, this matters for

pipeline design, error handling, and idempotency: multi-step loads/merges should group logically atomic DML steps into an explicit transaction so they commit together or roll back together, but DDL (like CREATE /ALTER) should be separated or executed with full awareness of implicit commits. This is consistent with SnowPro Architect expectations around reliable data engineering patterns and correctness under failure.

NEW QUESTION # 164

......