Latest Updated Snowflake Sample ARA-C01 Exam: SnowPro Advanced Architect Certification & ARA-C01 Exam Questions And Answers

BTW, DOWNLOAD part of Free4Torrent ARA-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1ZXT4dkooeDdLBChUfqbhwQe7vd2xgiac

Today is the right time to learn new and in demands skills. You can do this easily, just get registered in SnowPro Advanced Architect Certification ARA-C01 certification exam and start preparation with Snowflake ARA-C01 exam dumps. The SnowPro Advanced Architect Certification ARA-C01 pdf questions and practice test are ready for download. Just pay the affordable Snowflake ARA-C01 authentic dumps charges and click on the download button. Get the Channel Partner Program SnowPro Advanced Architect Certification ARA-C01 latest dumps and start preparing today.

To become certified in Snowflake ARA-C01, candidates must first meet the eligibility requirements, which include having experience in data warehousing, data modeling, and data integration. Candidates must also complete the SnowPro Core Certification Exam and pass the SnowPro Advanced Architect Certification Exam, which is a rigorous and comprehensive exam that tests advanced knowledge and skills in Snowflake architecture.

>> Sample ARA-C01 Exam <<

Eliminates confusion while taking the Snowflake ARA-C01 exam

Now on the Internet, a lot of online learning platform management is not standard, some web information may include some viruses,

cause far-reaching influence to pay end users and adverse effect. If you purchase our ARA-C01 test torrent this issue is impossible. We hire experienced staff to handle this issue perfectly. We are sure that our products and payment process are surely safe and anti-virus. If you have any question about downloading and using our ARA-C01 Study Tool, we have professional staff to remotely handle for you immediately, let users to use the SnowPro Advanced Architect Certification guide torrent in a safe environment, bring more comfortable experience for the user.

Snowflake SnowPro Advanced Architect Certification Sample Questions (Q71-Q76):

NEW QUESTION # 71
A company has a source system that provides JSON records for various loT operations. The JSON Is loading directly into a persistent table with a variant field. The data Is quickly growing to 100s of millions of records and performance to becoming an issue. There is a generic access pattern that Is used to filter on the create_date key within the variant field.
What can be done to improve performance?

- A. Alter the target table to include additional fields pulled from the JSON records. This would include a create_date field with a datatype of varchar. When this field is used in the filter, partition pruning will occur.
- B. Incorporate the use of multiple tables partitioned by date ranges. When a user or process needs to query a particular date range, ensure the appropriate base table Is used.
- C. Validate the size of the warehouse being used. If the record count is approaching 100s of millions, size XL will be the minimum size required to process this amount of data.
- D. Alter the target table to Include additional fields pulled from the JSON records. This would Include a create_date field with a datatype of time stamp. When this field Is used in the filter, partition pruning will occur.

Answer: D

Explanation:
* The correct answer is A because it improves the performance of queries by reducing the amount of data scanned and processed. By adding a create_date field with a timestamp data type, Snowflake can automatically cluster the table based on this field and prune the micro-partitions that do not match the filter condition. This avoids the need to parse the JSON data and access the variant field for every record.
* Option B is incorrect because it does not improve the performance of queries. By adding a create_date field with a varchar data type, Snowflake cannot automatically cluster the table based on this field and prune the micro-partitions that do not match the filter condition. This still requires parsing the JSON data and accessing the variant field for every record.
* Option C is incorrect because it does not address the root cause of the performance issue. By validating the size of the warehouse being used, Snowflake can adjust the compute resources to match the data volume and parallelize the query execution. However, this does not reduce the amount of data scanned and processed, which is the main bottleneck for queries on JSON data.
* Option D is incorrect because it adds unnecessary complexity and overhead to the data loading and querying process. By incorporating the use of multiple tables partitioned by date ranges, Snowflake can reduce the amount of data scanned and processed for queries that specify a date range. However, this requires creating and maintaining multiple tables, loading data into the appropriate table based on the date, and joining the tables for queries that span multiple date ranges. References:
* Snowflake Documentation: Loading Data Using Snowpipe: This document explains how to use Snowpipe to continuously load data from external sources into Snowflake tables. It also describes the syntax and usage of the COPY INTO command, which supports various options and parameters to control the loading behavior, such as ON_ERROR, PURGE, and SKIP_FILE.
* Snowflake Documentation: Date and Time Data Types and Functions: This document explains the different data types and functions for working with date and time values in Snowflake. It also describes how to set and change the session timezone and the system timezone.
* Snowflake Documentation: Querying Metadata: This document explains how to query the metadata of the objects and operations in Snowflake using various functions, views, and tables. It also describes how to access the copy history information using the COPY_HISTORY function or the COPY_HISTORY view.
* Snowflake Documentation: Loading JSON Data: This document explains how to load JSON data into Snowflake tables using various methods, such as the COPY INTO command, the INSERT command, or the PUT command. It also describes how to access and query JSON data using the dot notation, the FLATTEN function, or the LATERAL join.
* Snowflake Documentation: Optimizing Storage for Performance: This document explains how to optimize the storage of data in Snowflake tables to improve the performance of queries. It also describes the concepts and benefits of automatic clustering, search optimization service, and materialized views.

NEW QUESTION # 72
A healthcare company is deploying a Snowflake account that may include Personal Health Information (PHI). The company must

ensure compliance with all relevant privacy standards.

Which best practice recommendations will meet data protection and compliance requirements? (Choose three.)

- A. Use, at minimum, the Business Critical edition of Snowflake.
- B. Avoid sharing data with partner organizations.
- C. Create Dynamic Data Masking policies and apply them to columns that contain PHI.
- D. Rewrite SQL queries to eliminate projections of PHI data based on current_role().
- E. Use the External Tokenization feature to obfuscate sensitive data.
- F. Use the Internal Tokenization feature to obfuscate sensitive data.

Answer: A,C,E

NEW QUESTION # 73

The diagram shows the process flow for Snowpipe auto-ingest with Amazon Simple Notification Service (SNS) with the following steps:

Step 1: Data files are loaded in a stage.

Step 2: An Amazon S3 event notification, published by SNS, informs Snowpipe - by way of Amazon Simple Queue Service (SQS) - that files are ready to load. Snowpipe copies the files into a queue.

Step 3: A Snowflake-provided virtual warehouse loads data from the queued files into the target table based on parameters defined in the specified pipe.

If an AWS Administrator accidentally deletes the SQS subscription to the SNS topic in Step 2, what will happen to the pipe that references the topic to receive event messages from Amazon S3?

- A. The pipe will no longer be able to receive the messages and the user must wait for 24 hours from the time when the SNS topic subscription was deleted. Pipe recreation is not required as the pipe will reuse the same subscription to the existing SNS topic after 24 hours.
- B. The pipe will continue to receive the messages as Snowflake will automatically restore the subscription to the same SNS topic and will recreate the pipe by specifying the same SNS topic name in the pipe definition.
- C. The pipe will continue to receive the messages as Snowflake will automatically restore the subscription by creating a new SNS topic. Snowflake will then recreate the pipe by specifying the new SNS topic name in the pipe definition.
- D. The pipe will no longer be able to receive the messages. To restore the system immediately, the user needs to manually create a new SNS topic with a different name and then recreate the pipe by specifying the new SNS topic name in the pipe definition.

Answer: D

Explanation:

If an AWS Administrator accidentally deletes the SQS subscription to the SNS topic in Step 2, the pipe that references the topic to receive event messages from Amazon S3 will no longer be able to receive the messages. This is because the SQS subscription is the link between the SNS topic and the Snowpipe notification channel. Without the subscription, the SNS topic will not be able to send notifications to the Snowpipe queue, and the pipe will not be triggered to load the new files. To restore the system immediately, the user needs to manually create a new SNS topic with a different name and then recreate the pipe by specifying the new SNS topic name in the pipe definition. This will create a new notification channel and a new SQS subscription for the pipe. Alternatively, the user can also recreate the SQS subscription to the existing SNS topic and then alter the pipe to use the same SNS topic name in the pipe definition. This will also restore the notification channel and the pipe functionality. References:
Automating Snowpipe for Amazon S3
Enabling Snowpipe Error Notifications for Amazon SNS
HowTo: Configuration steps for Snowpipe Auto-Ingest with AWS S3 Stages
"To circumvent the 72-hour delay, you can create a SNS topic with a different name. Recreate any pipes that reference the topic using the CREATE OR REPLACE PIPE command, and specify the new topic name."
https://docs.snowflake.com/en/user-guide/data-load-snowpipe-ts#snowpipe-stops-loading-files-after-amazon- sns-topic-subscription-is-deleted

NEW QUESTION # 74

The following table exists in the production database:

A regulatory requirement states that the company must mask the username for events that are older than six months based on the current date when the data is queried.

How can the requirement be met without duplicating the event data and making sure it is applied when creating views using the table

or cloning the table?

- A. Use a masking policy on the username column with event_timestamp as a conditional column.
- B. Use a row level policy on the user_events table using a entitlement table with valid dates.
- C. Use a secure view on the user_events table using a case statement on the username column.
- D. Use a masking policy on the username column using a entitlement table with valid dates.

Answer: A

Explanation:
A masking policy is a feature of Snowflake that allows masking sensitive data in query results based on the role of the user and the condition of the data. A masking policy can be applied to a column in a table or a view, and it can use another column in the same table or view as a conditional column. A conditional column is a column that determines whether the masking policy is applied or not based on its value1.
In this case, the requirement can be met by using a masking policy on the username column with event_timestamp as a conditional column. The masking policy can use a function that masks the username if the event_timestamp is older than six months based on the current date, and returns the original username otherwise. The masking policy can be applied to the user_events table, and it will also be applied when creating views using the table or cloning the table2.
The other options are not correct because:
* A. Using a masking policy on the username column using an entitlement table with valid dates would require creating another table that stores the valid dates for each username, and joining it with the user_events table in the masking policy function. This would add complexity and overhead to the masking policy, and it would not use the event_timestamp column as the condition for masking.
* B. Using a row level policy on the user_events table using an entitlement table with valid dates would require creating another table that stores the valid dates for each username, and joining it with the user_events table in the row access policy function. This would filter out the rows that have event_timestamp older than six months based on the valid dates, instead of masking the username column. This would not meet the requirement of masking the username, and it would also reduce the visibility of the event data.
* D. Using a secure view on the user_events table using a case statement on the username column would require creating a view that uses a case expression to mask the username column based on the event_timestamp column. This would meet the requirement of masking the username, but it would not be applied when cloning the table. A secure view is a view that prevents the underlying data from being exposed by queries on the view. However, a secure view does not prevent the underlying data from being exposed by cloning the table3.
References:
* 1: Masking Policies | Snowflake Documentation
* 2: Using Conditional Columns in Masking Policies | Snowflake Documentation
* 3: Secure Views | Snowflake Documentation

NEW QUESTION # 75
If your role does not own the share, but owns the objects in the share, how can you block access to the objects

- A. Revoking the USAGE or SELECT privileges with CASCADE on the objects from the share owner.
- B. Revoking the USAGE or SELECT privileges on the objects from the share owner
- C. You will need to connect with the share owner

Answer: A

NEW QUESTION # 76
......

It is a common sense that only high quality and accuracy ARA-C01 practice materials can relive you from those worries. It is our communal wish to reap successful fruits. So our company did a lot to make sure that happen. Our ARA-C01 practice materials compiled by the most professional experts can offer you with high quality and accuracy results for your success. If you are unfamiliar with our ARA-C01 practice materials, please download the free demos for your reference, and to some unlearned exam candidates, you can master necessities by our ARA-C01 practice materials quickly.