

Latest ARA-C01 Exam Cost - ARA-C01 Free Exam Questions



P.S. Free 2026 Snowflake ARA-C01 dumps are available on Google Drive shared by TestKingFree:
<https://drive.google.com/open?id=1cJDmfhEPAGAdaYEMqD6jgoCoHg9A5CnI>

We hope you can feel that we sincerely hope to help you. We hope that after choosing our ARA-C01 study materials, you will be able to concentrate on learning our ARA-C01 learning guide without worry. It is our greatest honor that you can feel satisfied. Of course, we will value every user. We will never neglect any user. Our ARA-C01 Exam Braindumps will provide perfect service for everyone.

Snowflake ARA-C01 Exam covers a wide range of advanced topics related to Snowflake architecture, including data modeling, security, performance tuning, and data integration. Candidates are expected to have a deep understanding of Snowflake's features and functionality, as well as the ability to apply this knowledge to real-world scenarios.

>> Latest ARA-C01 Exam Cost <<

100% Pass Quiz 2026 Snowflake Unparalleled Latest ARA-C01 Exam Cost

Of course, the future is full of unknowns and challenges for everyone. Even so, we all hope that we can have a bright future. Pass the ARA-C01 exam, for most people, is an ability to live the life they want, and the realization of these goals needs to be established on a good basis of having a good job. A good job requires a certain amount of competence, and the most intuitive way to measure

competence is whether you get a series of the test ARA-C01 Certification and obtain enough qualifications.

Snowflake ARA-C01 (SnowPro Advanced Architect Certification) Certification Exam is a highly reputable certification that is recognized globally by businesses and organizations that use Snowflake. SnowPro Advanced Architect Certification certification exam is designed to test the skills and knowledge of individuals who want to become advanced architects in data warehousing and data analytics. SnowPro Advanced Architect Certification certification is a valuable asset for individuals who want to advance their careers in these fields, and there are several resources available to help candidates prepare for the exam.

Snowflake SnowPro Advanced Architect Certification Sample Questions (Q153-Q158):

NEW QUESTION # 153

Arrange in order of performance(least to high)

1. External Table
2. External Table with partitioning
3. Materialized View Vs External table
4. Directly querying a file from the stage

- A. 3,2,4,1
- B. 1,2,4,3
- C. 4,3,2,1
- D. 1,2,3,4

Answer: D

NEW QUESTION # 154

An Architect needs to design a data unloading strategy for Snowflake, that will be used with the COPY INTO <location> command.

Which configuration is valid?

- A. Location of files: Azure ADLS. File formats: JSON, XML, Avro, Parquet, ORC. Compression: bzip2. Encryption: User-supplied key
- B. Location of files: Snowflake internal location. File formats: CSV, XML. File encoding: UTF-8. Encryption: 128-bit
- C. Location of files: Amazon S3. File formats: CSV, JSON. File encoding: Latin-1 (ISO-8859). Encryption: 128-bit
- D. Location of files: Google Cloud Storage. File formats: Parquet. File encoding: UTF-8 Compression: gzip

Answer: D

Explanation:

For the configuration of data unloading in Snowflake, the valid option among the provided choices is "C." This is because Snowflake supports unloading data into Google Cloud Storage using the COPY INTO <location> command with specific configurations. The configurations listed in option C, such as Parquet file format with UTF-8 encoding and gzip compression, are all supported by Snowflake. Notably, Parquet is a columnar storage file format, which is optimal for high-performance data processing tasks in Snowflake. The UTF-8 file encoding and gzip compression are both standard and widely used settings that are compatible with Snowflake's capabilities for data unloading to cloud storage platforms.

References:

[Snowflake Documentation on COPY INTO command](#)

[Snowflake Documentation on Supported File Formats](#)

[Snowflake Documentation on Compression and Encoding Options](#)

NEW QUESTION # 155

What is the MOST efficient way to design an environment where data retention is not considered critical, and customization needs are to be kept to a minimum?

- A. Use a temporary table.

- B. Use a transient table.
- **C. Use a transient database.**
- D. Use a transient schema.

Answer: C

Explanation:

Transient databases in Snowflake are designed for situations where data retention is not critical, and they do not have the fail-safe period that regular databases have. This means that data in a transient database is not recoverable after the Time Travel retention period. Using a transient database is efficient because it minimizes storage costs while still providing most functionalities of a standard database without the overhead of data protection features that are not needed when data retention is not a concern.

NEW QUESTION # 156

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

Answer: C

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.

1: Masking Policies | Snowflake Documentation

2: Using Conditional Columns in Masking Policies | Snowflake Documentation

3: Secure Views | Snowflake Documentation

NEW QUESTION # 157

A user is executing the following command sequentially within a timeframe of 10 minutes from start to finish:

What would be the output of this query?

- A. Time Travel data is not available for table T_SALES.
- B. The offset -> is not a valid clause in the clone operation.
- C. Syntax error line 1 at position 58 unexpected 'at'.

- D. Table T_SALES_CLONE successfully created.

Answer: D

Explanation:

The query is executing a clone operation on an existing table `t_sales` with an offset to account for the retention time. The syntax used is correct for cloning a table in Snowflake, and the use of the `at(offset => -60*30)` clause is valid. This specifies that the clone should be based on the state of the table 30 minutes prior (60 seconds * 30). Assuming the table `t_sales` exists and has been modified within the last 30 minutes, and considering the `data_retention_time_in_days` is set to 1 day (which enables time travel queries for the past

24 hours), the table `t_sales_clone` would be successfully created based on the state of `t_sales` 30 minutes before the clone command was issued.

NEW QUESTION # 158

• • • • •

ARA-C01 Free Exam Questions: <https://www.testkingfree.com/Snowflake/ARA-C01-practice-exam-dumps.html>

BTW, DOWNLOAD part of TestKingFree ARA-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=1cJDmfhEPAGAdaYEMqD6jgoCoHg9A5CnI>