

ARA-C01^{Q&As}

SnowPro Advanced: Architect Certification Exam

Pass Snowflake ARA-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass2lead.com/ara-c01.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Snowflake
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



QUESTION 1

A company has a table with that has corrupted data, named Data. The company wants to recover the data as it was 5 minutes ago using cloning and Time Travel.

What command will accomplish this?

- A. `CREATE CLONE TABLE Recover_Data FROM Data AT(OFFSET => -60*5);`
- B. `CREATE CLONE Recover_Data FROM Data AT(OFFSET => -60*5);`
- C. `CREATE TABLE Recover_Data CLONE Data AT(OFFSET => -60*5);`
- D. `CREATE TABLE Recover Data CLONE Data AT(TIME => -60*5);`

Correct Answer: C

Explanation: This is the correct command to create a clone of the table Data as it was 5 minutes ago using cloning and Time Travel. Cloning is a feature that allows creating a copy of a database, schema, table, or view without duplicating the

data or metadata. Time Travel is a feature that enables accessing historical data (i.e. data that has been changed or deleted) at any point within a defined period. To create a clone of a table at a point in time in the past, the syntax is:

```
CREATE TABLE CLONE AT (OFFSET => );
```

The OFFSET parameter specifies the time difference in seconds from the present time. A negative value indicates a point in the past. For example, -60*5 means 5 minutes ago. Alternatively, the TIMESTAMP parameter can be used to specify

an exact timestamp in the past. The clone will contain the data as it existed in the source table at the specified point in time.

References:

Snowflake Documentation: Cloning Objects

Snowflake Documentation: Cloning Objects at a Point in Time in the Past

QUESTION 2

How does a standard virtual warehouse policy work in Snowflake?

- A. It conserves credits by keeping running clusters fully loaded rather than starting additional clusters.
- B. It starts only if the system estimates that there is a query load that will keep the cluster busy for at least 6 minutes.
- C. It starts only if the system estimates that there is a query load that will keep the cluster busy for at least 2 minutes.
- D. It prevents or minimizes queuing by starting additional clusters instead of conserving credits.

Correct Answer: D

Explanation: A standard virtual warehouse policy is one of the two scaling policies available for multi-cluster warehouses in Snowflake. The other policy is economic. A standard policy aims to prevent or minimize queuing by starting additional clusters as soon as the current cluster is fully loaded, regardless of the number of queries in the queue. This policy can improve query performance and concurrency, but it may also consume more credits than an economic policy, which tries to conserve credits by keeping the running clusters fully loaded before starting additional clusters. The scaling policy can be set when creating or modifying a warehouse, and it can be changed at any time. References: Snowflake Documentation: Multi-cluster Warehouses Snowflake Documentation: Scaling Policy for Multi-cluster Warehouses

QUESTION 3

What Snowflake features should be leveraged when modeling using Data Vault? (Choose two.)

- A. Snowflake's support of multi-table inserts into the data model's Data Vault tables
- B. Data needs to be pre-partitioned to obtain a superior data access performance
- C. Scaling up the virtual warehouses will support parallel processing of new source loads
- D. Snowflake's ability to hash keys so that hash key joins can run faster than integer joins

Correct Answer: AC

Explanation: These two features are relevant for modeling using Data Vault on Snowflake. Data Vault is a data modeling approach that organizes data into hubs, links, and satellites. Data Vault is designed to enable high scalability, flexibility, and performance for data integration and analytics. Snowflake is a cloud data platform that supports various data modeling techniques, including Data Vault. Snowflake provides some features that can enhance the Data Vault modeling, such as: Snowflake's support of multi-table inserts into the data model's Data Vault tables. Multi-table inserts (MTI) are a feature that allows inserting data from a single query into multiple tables in a single DML statement. MTI can improve the performance and efficiency of loading data into Data Vault tables, especially for real-time or near-real-time data integration. MTI can also reduce the complexity and maintenance of the loading code, as well as the data duplication and latency¹². Scaling up the virtual warehouses will support parallel processing of new source loads. Virtual warehouses are a feature that allows provisioning compute resources on demand for data processing. Virtual warehouses can be scaled up or down by changing the size of the warehouse, which determines the number of servers in the warehouse. Scaling up the virtual warehouses can improve the performance and concurrency of processing new source loads into Data Vault tables, especially for large or complex data sets. Scaling up the virtual warehouses can also leverage the parallelism and distribution of Snowflake's architecture, which can optimize the data loading and querying³⁴. References: Snowflake Documentation: Multi-table Inserts Snowflake Blog: Tips for Optimizing the Data Vault Architecture on Snowflake Snowflake Documentation: Virtual Warehouses Snowflake Blog: Building a Real-Time Data Vault in Snowflake

QUESTION 4

At which object type level can the APPLY MASKING POLICY, APPLY ROW ACCESS POLICY and APPLY SESSION POLICY privileges be granted?

- A. Global
- B. Database
- C. Schema
- D. Table

Correct Answer: A

Explanation: The object type level at which the APPLY MASKING POLICY, APPLY ROW ACCESS POLICY and APPLY SESSION POLICY privileges can be granted is global. These are account-level privileges that control who can apply or unset these policies on objects such as columns, tables, views, accounts, or users. These privileges are granted to the ACCOUNTADMIN role by default, and can be granted to other roles as needed. The other options are incorrect because they are not the object type level at which these privileges can be granted. Database, schema, and table are lower-level object types that do not support these privileges. References: Access Control Privileges | Snowflake Documentation, Using Dynamic Data Masking | Snowflake Documentation, Using Row Access Policies | Snowflake Documentation, Using Session Policies | Snowflake Documentation

QUESTION 5

An Architect would like to save quarter-end financial results for the previous six years.

Which Snowflake feature can the Architect use to accomplish this?

- A. Search optimization service
- B. Materialized view
- C. Time Travel
- D. Zero-copy cloning
- E. Secure views

Correct Answer: D

Explanation: Zero-copy cloning is a Snowflake feature that can be used to save quarter-end financial results for the previous six years. Zero-copy cloning allows creating a copy of a database, schema, table, or view without duplicating the data or metadata. The clone shares the same data files as the original object, but tracks any changes made to the clone or the original separately. Zero-copy cloning can be used to create snapshots of data at different points in time, such as quarter-end financial results, and preserve them for future analysis or comparison. Zero-copy cloning is fast, efficient, and does not consume any additional storage space unless the data is modified¹. References: Zero-Copy Cloning | Snowflake Documentation

[ARA-C01 PDF Dumps](#)

[ARA-C01 VCE Dumps](#)

[ARA-C01 Exam Questions](#)