

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

A healthcare company wants to share data with a medical institute. The institute is running a Standard edition of Snowflake; the healthcare company is running a Business Critical edition.

How can this data be shared?

- A. The healthcare company will need to change the institute's Snowflake edition in the accounts panel.
- B. By default, sharing is supported from a Business Critical Snowflake edition to a Standard edition.
- C. Contact Snowflake and they will execute the share request for the healthcare company.
- D. Set the share_restriction parameter on the shared object to false.

Correct Answer: D

Explanation: By default, Snowflake does not allow sharing data from a Business Critical edition to a non-Business Critical edition. This is because Business Critical edition provides enhanced security and data protection features that are not available in lower editions. However, this restriction can be overridden by setting the `share_restriction` parameter on the shared object (database, schema, or table) to false. This parameter allows the data provider to explicitly allow sharing data with lower edition accounts. Note that this parameter can only be set by the data provider, not the data consumer. Also, setting this parameter to false may reduce the level of security and data protection for the shared data. References: Enable Data Share: Business Critical Account to Lower Edition Sharing Is Not Allowed From An Account on BUSINESS CRITICAL Edition to an Account On A Lower Edition SQL Execution Error: Sharing is Not Allowed from an Account on BUSINESS CRITICAL Edition to an Account on a Lower Edition Snowflake Editions | Snowflake Documentation

QUESTION 3

How is the change of local time due to daylight savings time handled in Snowflake tasks? (Choose two.)

- A. A task scheduled in a UTC-based schedule will have no issues with the time changes.
- B. Task schedules can be designed to follow specified or local time zones to accommodate the time changes.
- C. A task will move to a suspended state during the daylight savings time change.
- D. A frequent task execution schedule like minutes may not cause a problem, but will affect the task history.
- E. A task schedule will follow only the specified time and will fail to handle lost or duplicated hours.

Correct Answer: AB

Explanation: According to the Snowflake documentation¹ and the web search results², these two statements are true about how the change of local time due to daylight savings time is handled in Snowflake tasks. A task is a feature that allows scheduling and executing SQL statements or stored procedures in Snowflake. A task can be scheduled using a cron expression that specifies the frequency and time zone of the task execution. A task scheduled in a UTC-based schedule will have no issues with the time changes. UTC is a universal time standard that does not observe daylight savings time. Therefore, a task that uses UTC as the time zone will run at the same time throughout the year, regardless of the local time changes¹. Task schedules can be designed to follow specified or local time zones to accommodate the time changes. Snowflake supports using any valid IANA time zone identifier in the cron expression for a task. This allows the task to run according to the local time of the specified time zone, which may include daylight savings time adjustments. For example, a task that uses Europe/London as the time zone will run one hour earlier or later when the local time switches between GMT and BST¹². References: Snowflake Documentation: Scheduling Tasks Snowflake Community: Do the timezones used in scheduling tasks in Snowflake adhere to daylight savings?

QUESTION 4

Company A would like to share data in Snowflake with Company B. Company B is not on the same cloud platform as Company A.

What is required to allow data sharing between these two companies?

- A. Create a pipeline to write shared data to a cloud storage location in the target cloud provider.
- B. Ensure that all views are persisted, as views cannot be shared across cloud platforms.

- C. Setup data replication to the region and cloud platform where the consumer resides.
- D. Company A and Company B must agree to use a single cloud platform: Data sharing is only possible if the companies share the same cloud provider.

Correct Answer: C

Explanation: According to the SnowPro Advanced: Architect documents and learning resources, the requirement to allow data sharing between two companies that are not on the same cloud platform is to set up data replication to the region and cloud platform where the consumer resides. Data replication is a feature of Snowflake that enables copying databases across accounts in different regions and cloud platforms. Data replication allows data providers to securely share data with data consumers across different regions and cloud platforms by creating a replica database in the consumer's account. The replica database is read-only and automatically synchronized with the primary database in the provider's account. Data replication is useful for scenarios where data sharing is not possible or desirable due to latency, compliance, or security reasons¹. The other options are incorrect because they are not required or feasible to allow data sharing between two companies that are not on the same cloud platform. Option A is incorrect because creating a pipeline to write shared data to a cloud storage location in the target cloud provider is not a secure or efficient way of sharing data. It would require additional steps to load the data from the cloud storage to the consumer's account, and it would not leverage the benefits of Snowflake's data sharing features. Option B is incorrect because ensuring that all views are persisted is not relevant for data sharing across cloud platforms. Views can be shared across cloud platforms as long as they reference objects in the same database. Persisting views is an option to improve the performance of querying views, but it is not required for data sharing². Option D is incorrect because Company A and Company B do not need to agree to use a single cloud platform. Data sharing is possible across different cloud platforms using data replication or other methods, such as listings or auto- fulfillment³. References: Replicating Databases Across Multiple Accounts | Snowflake Documentation, Persisting Views | Snowflake Documentation, Sharing Data Across Regions and Cloud Platforms | Snowflake Documentation

QUESTION 5

An Architect is designing a pipeline to stream event data into Snowflake using the Snowflake Kafka connector. The Architect's highest priority is to configure the connector to stream data in the MOST cost-effective manner.

Which of the following is recommended for optimizing the cost associated with the Snowflake Kafka connector?

- A. Utilize a higher Buffer.flush.time in the connector configuration.
- B. Utilize a higher Buffer.size.bytes in the connector configuration.
- C. Utilize a lower Buffer.size.bytes in the connector configuration.
- D. Utilize a lower Buffer.count.records in the connector configuration.

Correct Answer: A

Explanation: The minimum value supported for the buffer.flush.time property is 1 (in seconds). For higher average data flow rates, we suggest that you decrease the default value for improved latency. If cost is a greater concern than latency, you could increase the buffer flush time. Be careful to flush the Kafka memory buffer before it becomes full to avoid out of memory exceptions. <https://docs.snowflake.com/en/user-guide/data-load-snowpipe-streaming-kafka>