



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.pass4itsure.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**

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

Correct Answer: AC

Explanation: These two options are the best ways to meet the requirement of loading data from an external stage and making it accessible by dashboards with the least amount of coding. Snowpipe with auto-ingest is a feature that enables continuous and automated data loading from an external stage into a Snowflake table. Snowpipe uses event notifications from the cloud storage service to detect new or modified files in the stage and triggers a COPY INTO command to load the data into the table. Snowpipe is efficient, scalable, and serverless, meaning it does not require any infrastructure or maintenance from the user. Snowpipe also supports loading data from files of any size, as long as they are in a supported format¹. A materialized view on an external table is a feature that enables creating a pre-computed result set from an external table and storing it in Snowflake. A materialized view can improve the performance and efficiency of querying data from an external table, especially for complex queries or dashboards. A materialized view can also support aggregations, joins, and filters on the external table data. A materialized view on an external table is automatically refreshed when the underlying data in the external stage changes, as long as the AUTO_REFRESH parameter is set to true². References: Snowpipe Overview | Snowflake Documentation Materialized Views on External Tables | Snowflake Documentation

QUESTION 2

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 3

Following objects can be cloned in snowflake: (Choose three.)

- A. Permanent table
- B. Transient table
- C. Temporary table
- D. External tables
- E. Internal stages

Correct Answer: ABD

Snowflake supports cloning of various objects, such as databases, schemas, tables, stages, file formats, sequences, streams, tasks, and roles. Cloning creates a copy of an existing object in the system without copying the data or metadata.

Cloning is also known as zero-copy cloning¹. Among the objects listed in the question, the following ones can be cloned in Snowflake:

The following objects listed in the question cannot be cloned in Snowflake:

References: : Cloning Considerations : CREATE TABLE ... CLONE : CREATE EXTERNAL TABLE ... CLONE : Temporary Tables : Internal Stages

QUESTION 4

There are two databases in an account, named fin_db and hr_db which contain payroll and employee data, respectively. Accountants and Analysts in the company require different permissions on the objects in these databases to perform their jobs. Accountants need read-write access to fin_db but only require read-only access to hr_db because the database is maintained by human resources personnel.

An Architect needs to create a read-only role for certain employees working in the human resources department.

Which permission sets must be granted to this role?

- A. USAGE on database hr_db, USAGE on all schemas in database hr_db, SELECT on all tables in database hr_db
- B. USAGE on database hr_db, SELECT on all schemas in database hr_db, SELECT on all tables in database hr_db
- C. MODIFY on database hr_db, USAGE on all schemas in database hr_db, USAGE on all tables in database hr_db



D. USAGE on database hr_db, USAGE on all schemas in database hr_db, REFERENCES on all tables in database hr_db

Correct Answer: A

To create a read-only role for certain employees working in the human resources department, the role needs to have the following permissions on the hr_db database: Option A is the correct answer because it grants the minimum permissions required for a read-only role on the hr_db database. Option B is incorrect because SELECT on schemas is not a valid permission. Schemas only support USAGE and CREATE permissions. Option C is incorrect because MODIFY on the database is not a valid permission. Databases only support USAGE, CREATE, MONITOR, and OWNERSHIP permissions. Moreover, USAGE on tables is not sufficient for querying the data. Tables support SELECT, INSERT, UPDATE, DELETE, TRUNCATE, REFERENCES, and OWNERSHIP permissions. Option D is incorrect because REFERENCES on tables is not relevant for querying the data. REFERENCES permission allows the role to create foreign key constraints on the tables. References: : <https://docs.snowflake.com/en/user-guide/security-access-control-privileges.html#database-privileges> : <https://docs.snowflake.com/en/user-guide/security-access-control-privileges.html#schema-privileges> : <https://docs.snowflake.com/en/user-guide/security-access-control-privileges.html#table-privileges>

QUESTION 5

A company has a Snowflake account named ACCOUNTA in AWS us-east-1 region. The company stores its marketing data in a Snowflake database named MARKET_DB. One of the company's business partners has an account named PARTNERB in Azure East US 2 region. For marketing purposes the company has agreed to share the database MARKET_DB with the partner account.

Which of the following steps MUST be performed for the account PARTNERB to consume data from the MARKET_DB database?

- A. Create a new account (called AZABC123) in Azure East US 2 region. From account ACCOUNTA create a share of database MARKET_DB, create a new database out of this share locally in AWS us-east-1 region, and replicate this new database to AZABC123 account. Then set up data sharing to the PARTNERB account.
- B. From account ACCOUNTA create a share of database MARKET_DB, and create a new database out of this share locally in AWS us-east-1 region. Then make this database the provider and share it with the PARTNERB account.
- C. Create a new account (called AZABC123) in Azure East US 2 region. From account ACCOUNTA replicate the database MARKET_DB to AZABC123 and from this account set up the data sharing to the PARTNERB account.
- D. Create a share of database MARKET_DB, and create a new database out of this share locally in AWS us-east-1 region. Then replicate this database to the partner's account PARTNERB.

Correct Answer: C

Snowflake supports data sharing across regions and cloud platforms using account replication and share replication features. Account replication enables the replication of objects from a source account to one or more target accounts in the same organization. Share replication enables the replication of shares from a source account to one or more target accounts in the same organization¹. To share data from the MARKET_DB database in the ACCOUNTA account in AWS useast-1 region with the PARTNERB account in Azure East US 2 region, the following steps must be performed: Therefore, option C is the correct answer. References: : [Replicating Shares Across Regions and Cloud Platforms](#) : [Working with Organizations and Accounts](#) : [Replicating Databases Across Multiple Accounts](#) : [Replicating Shares Across Multiple Accounts](#)