

100% Money Back
Guarantee

Vendor:Snowflake

Exam Code:ARA-C01

Exam Name:SnowPro Advanced: Architect
Certification Exam

Version:Demo

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

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

The following DDL command was used to create a task based on a stream:

```
CREATE TASK ts_insert_new_customers
  WAREHOUSE = MY_WH
  Schedule = '5 minute'
WHEN
  System$STREAM_HAS_DATA('MYSTREAM')
AS
  INSERT INTO new_customers(id, name) SELECT id, name
  FROM mystream WHERE METADATA$action = 'INSERT';
```

Assuming MY_WH is set to auto_suspend = 60 and used exclusively for this task, which statement is true?

- A. The warehouse MY_WH will be made active every five minutes to check the stream.
- B. The warehouse MY_WH will only be active when there are results in the stream.
- C. The warehouse MY_WH will never suspend.
- D. The warehouse MY_WH will automatically resize to accommodate the size of the stream.

Correct Answer: B

Explanation: The warehouse MY_WH will only be active when there are results in the stream. This is because the task is created based on a stream, which means that the task will only be executed when there are new data in the stream.

Additionally, the warehouse is set to auto_suspend = 60, which means that the warehouse will automatically suspend after 60 seconds of inactivity. Therefore, the warehouse will only be active when there are results in the stream.

References:

[CREATE TASK | Snowflake Documentation]

[Using Streams and Tasks | Snowflake Documentation] [CREATE WAREHOUSE | Snowflake Documentation]

QUESTION 4

Which of the following are characteristics of Snowflake's parameter hierarchy?

- A. Session parameters override virtual warehouse parameters.
- B. Virtual warehouse parameters override user parameters.
- C. Table parameters override virtual warehouse parameters.
- D. Schema parameters override account parameters.

Correct Answer: D

Explanation: This is the correct answer because it reflects the characteristics of Snowflake's parameter hierarchy. Snowflake provides three types of parameters that can be set for an account: account parameters, session parameters, and object parameters. All parameters have default values, which can be set and then overridden at different levels depending on the parameter type. The following diagram illustrates the hierarchical relationship between the different parameter types and how individual parameters can be overridden at each level¹: As shown in the diagram, schema parameters are a type of object parameters that can be set for schemas. Schema parameters can override the account parameters that are set at the account level. For example, the LOG_LEVEL parameter can be set at the account level to control the logging level for all objects in the account, but it can also be overridden at the schema level to control the logging level for specific stored procedures and UDFs in that schema². The other options listed are not correct because they do not reflect the characteristics of Snowflake's parameter hierarchy. Session parameters do not override virtual warehouse parameters, because virtual warehouse parameters are a type of session parameters that can be set for virtual warehouses. Virtual warehouse parameters do not override user parameters, because user parameters are a type of session parameters that can be set for users. Table parameters do not override virtual warehouse parameters, because table parameters are a type of object parameters that can be set for tables, and object parameters do not affect session parameters¹. References: Snowflake Documentation: Parameters Snowflake Documentation: Setting Log Level

QUESTION 5

A DevOps team has a requirement for recovery of staging tables used in a complex set of data pipelines. The staging tables are all located in the same staging schema. One of the requirements is to have online recovery of data on a rolling 7day basis.

After setting up the DATA_RETENTION_TIME_IN_DAYS at the database level, certain tables remain unrecoverable past 1 day.

What would cause this to occur? (Choose two.)

- A. The staging schema has not been setup for MANAGED ACCESS.
- B. The DATA_RETENTION_TIME_IN_DAYS for the staging schema has been set to 1 day.
- C. The tables exceed the 1 TB limit for data recovery.
- D. The staging tables are of the TRANSIENT type.
- E. The DevOps role should be granted ALLOW_RECOVERY privilege on the staging schema.

Correct Answer: BD

The DATA_RETENTION_TIME_IN_DAYS parameter controls the Time Travel retention period for an object (database, schema, or table) in Snowflake. This parameter specifies the number of days for which historical data is preserved and can be accessed using Time Travel operations (SELECT, CREATE ... CLONE, UNDROP)¹. The requirement for

recovery of staging tables on a rolling 7-day basis means that the `DATA_RETENTION_TIME_IN_DAYS` parameter should be set to 7 at the database level. However, this parameter can be overridden at the lower levels (schema or table) if they have a different value¹. Therefore, one possible cause for certain tables to remain unrecoverable past 1 day is that the `DATA_RETENTION_TIME_IN_DAYS` for the staging schema has been set to 1 day. This would override the database level setting and limit the Time Travel retention period for all the tables in the schema to 1 day. To fix this, the parameter should be unset or set to 7 at the schema level¹. Therefore, option B is correct. Another possible cause for certain tables to remain unrecoverable past 1 day is that the staging tables are of the `TRANSIENT` type. Transient tables are tables that do not have a Fail-safe period and can have a Time Travel retention period of either 0 or 1 day. Transient tables are suitable for temporary or intermediate data that can be easily reproduced or replicated². To fix this, the tables should be created as permanent tables, which can have a Time Travel retention period of up to 90 days¹. Therefore, option D is correct. Option A is incorrect because the `MANAGED ACCESS` feature is not related to the data recovery requirement. `MANAGED ACCESS` is a feature that allows granting access privileges to objects without explicitly granting the privileges to roles. It does not affect the Time Travel retention period or the data availability³. Option C is incorrect because there is no 1 TB limit for data recovery in Snowflake. The data storage size does not affect the Time Travel retention period or the data availability⁴. Option E is incorrect because there is no `ALLOW_RECOVERY` privilege in Snowflake. The privilege required to perform Time Travel operations is `SELECT`, which allows querying historical data in tables⁵. References: : Understanding and Using Time Travel : Transient Tables : Managed Access : Understanding Storage Cost : Table Privileges

QUESTION 6

An Architect runs the following SQL query:

```
SELECT
  METADATA$FILENAME,
  METADATA$FILE_ROW_NUMBER
FROM @FILEROWS/Food_Reviews.csv
   (file_format=CSV_N)
```

How can this query be interpreted?

- A. `FILEROWS` is a stage. `FILE_ROW_NUMBER` is line number in file.
- B. `FILEROWS` is the table. `FILE_ROW_NUMBER` is the line number in the table.
- C. `FILEROWS` is a file. `FILE_ROW_NUMBER` is the file format location.
- D. `FILERONS` is the file format location. `FILE_ROW_NUMBER` is a stage.

Correct Answer: A

A stage is a named location in Snowflake that can store files for data loading and unloading. A stage can be internal or external, depending on where the files are stored.

The query in the question uses the `LIST` function to list the files in a stage named `FILEROWS`. The function returns a table with various columns, including `FILE_ROW_NUMBER`, which is the line number of the file in the stage. Therefore, the

query can be interpreted as listing the files in a stage named FILEROWS and showing the line number of each file in the stage.

References:

: Stages

: LIST Function

QUESTION 7

A user can change object parameters using which of the following roles?

- A. ACCOUNTADMIN, SECURITYADMIN
- B. SYSADMIN, SECURITYADMIN
- C. ACCOUNTADMIN, USER with PRIVILEGE
- D. SECURITYADMIN, USER with PRIVILEGE

Correct Answer: C

Explanation: According to the Snowflake documentation, object parameters are parameters that can be set on individual objects such as databases, schemas, tables, and stages. Object parameters can be set by users with the appropriate privileges on the objects. For example, to set the object parameter AUTO_REFRESH on a table, the user must have the MODIFY privilege on the table. The ACCOUNTADMIN role has the highest level of privileges on all objects in the account, so it can set any object parameter on any object. However, other roles, such as SECURITYADMIN or SYSADMIN, do not have the same level of privileges on all objects, so they cannot set object parameters on objects they do not own or have the required privileges on. Therefore, the correct answer is C. ACCOUNTADMIN, USER with PRIVILEGE. References: Parameters | Snowflake Documentation Object Parameters | Snowflake Documentation Object Privileges | Snowflake Documentation

QUESTION 8

A Data Engineer is designing a near real-time ingestion pipeline for a retail company to ingest event logs into Snowflake to derive insights. A Snowflake Architect is asked to define security best practices to configure access control privileges for the data load for auto- ingest to Snowpipe.

What are the MINIMUM object privileges required for the Snowpipe user to execute Snowpipe?

- A. OWNERSHIP on the named pipe, USAGE on the named stage, target database, and schema, and INSERT and SELECT on the target table
- B. OWNERSHIP on the named pipe, USAGE and READ on the named stage, USAGE on the target database and schema, and INSERT and SELECT on the target table
- C. CREATE on the named pipe, USAGE and READ on the named stage, USAGE on the target database and schema, and INSERT and SELECT on the target table
- D. USAGE on the named pipe, named stage, target database, and schema, and INSERT and SELECT on the target table

Correct Answer: B

Explanation: According to the SnowPro Advanced: Architect documents and learning resources, the minimum object privileges required for the Snowpipe user to execute Snowpipe are:

OWNERSHIP on the named pipe. This privilege allows the Snowpipe user to create, modify, and drop the pipe object that defines the COPY statement for loading data from the stage to the table¹.

USAGE and READ on the named stage. These privileges allow the Snowpipe user to access and read the data files from the stage that are loaded by Snowpipe². USAGE on the target database and schema. These privileges allow the

Snowpipe user to access the database and schema that contain the target table³. INSERT and SELECT on the target table. These privileges allow the Snowpipe user to insert data into the table and select data from the table⁴. The other

options are incorrect because they do not specify the minimum object privileges required for the Snowpipe user to execute Snowpipe. Option A is incorrect because it does not include the READ privilege on the named stage, which is required

for the Snowpipe user to read the data files from the stage. Option C is incorrect because it does not include the OWNERSHIP privilege on the named pipe, which is required for the Snowpipe user to create, modify, and drop the pipe object.

Option D is incorrect because it does not include the OWNERSHIP privilege on the named pipe or the READ privilege on the named stage, which are both required for the Snowpipe user to execute Snowpipe. References : CREATE PIPE |

Snowflake Documentation, CREATE STAGE | Snowflake Documentation, CREATE DATABASE | Snowflake Documentation, CREATE TABLE | Snowflake Documentation

QUESTION 9

A company has several sites in different regions from which the company wants to ingest data.

Which of the following will enable this type of data ingestion?

- A. The company must have a Snowflake account in each cloud region to be able to ingest data to that account.
- B. The company must replicate data between Snowflake accounts.
- C. The company should provision a reader account to each site and ingest the data through the reader accounts.
- D. The company should use a storage integration for the external stage.

Correct Answer: D

Explanation: This is the correct answer because it allows the company to ingest data from different regions using a storage integration for the external stage. A storage integration is a feature that enables secure and easy access to files in external cloud storage from Snowflake. A storage integration can be used to create an external stage, which is a named location that references the files in the external storage. An external stage can be used to load data into Snowflake tables using the COPY INTO command, or to unload data from Snowflake tables using the COPY INTO LOCATION command. A storage integration can support multiple regions and cloud platforms, as long as the external storage service is compatible with Snowflake¹². References: Snowflake Documentation: Storage Integrations
Snowflake Documentation: External Stages

QUESTION 10

Which of the below commands will use warehouse credits? (Choose three.)

- A. SHOW TABLES LIKE '\\SNOWFL%\\';
- B. SELECT MAX(FLAKE_ID) FROM SNOWFLAKE;
- C. SELECT COUNT(*) FROM SNOWFLAKE;
- D. SELECT COUNT(FLAKE_ID) FROM SNOWFLAKE GROUP BY FLAKE_ID;

Correct Answer: BCD

Warehouse credits are used to pay for the processing time used by each virtual warehouse in Snowflake. A virtual warehouse is a cluster of compute resources that enables executing queries, loading data, and performing other DML

operations. Warehouse credits are charged based on the number of virtual warehouses you use, how long they run, and their size¹. Among the commands listed in the question, the following ones will use warehouse credits:

The command that will not use warehouse credits is:

References: : Understanding Compute Cost : MAX Function : COUNT Function : GROUP BY Clause : SHOW TABLES

QUESTION 11

An Architect entered the following commands in sequence:

```
CREATE DATABASE SANDBOX;  
CREATE ROLE INTERN;  
CREATE TABLE SANDBOX.PUBLIC.AGENDA (ID INT, ITEMS STRING);  
GRANT SELECT ON ALL TABLES IN SCHEMA SANDBOX.PUBLIC TO ROLE INTERN;  
GRANT ROLE INTERN TO USER USER1;
```

USER1 cannot find the table.

Which of the following commands does the Architect need to run for USER1 to find the tables using the Principle of Least Privilege? (Choose two.)

- A. GRANT ROLE PUBLIC TO ROLE INTERN;
- B. GRANT USAGE ON DATABASE SANDBOX TO ROLE INTERN;
- C. GRANT USAGE ON SCHEMA SANDBOX.PUBLIC TO ROLE INTERN;
- D. GRANT OWNERSHIP ON DATABASE SANDBOX TO USER INTERN;
- E. GRANT ALL PRIVILEGES ON DATABASE SANDBOX TO ROLE INTERN;

Correct Answer: BC

According to the Principle of Least Privilege, the Architect should grant the minimum privileges necessary for the USER1 to find the tables in the SANDBOX database.

The USER1 needs to have USAGE privilege on the SANDBOX database and the SANDBOX.PUBLIC schema to be able to access the tables in the PUBLIC schema. Therefore, the commands B and C are the correct ones to run. The

command A is not correct because the PUBLIC role is automatically granted to every user and role in the account, and it does not have any privileges on the SANDBOX database by default.

The command D is not correct because it would transfer the ownership of the SANDBOX database from the Architect to the USER1, which is not necessary and violates the Principle of Least Privilege.

The command E is not correct because it would grant all the possible privileges on the SANDBOX database to the USER1, which is also not necessary and violates the Principle of Least Privilege.

References: : Snowflake - Principle of Least Privilege : Snowflake - Access Control Privileges : Snowflake - Public Role : Snowflake - Ownership and Grants

QUESTION 12

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