



# ARA-C01<sup>Q&As</sup>

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

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 2

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 3

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 4

An Architect has chosen to separate their Snowflake Production and QA environments using two separate Snowflake accounts.

The QA account is intended to run and test changes on data and database objects before pushing those changes to the Production account. It is a requirement that all database objects and data in the QA account need to be an exact copy of the database objects, including privileges and data in the Production account on at least a nightly basis.

Which is the LEAST complex approach to use to populate the QA account with the Production account's data and database objects on a nightly basis?

- A. 1) Create a share in the Production account for each database 2) Share access to the QA account as a Consumer 3) The QA account creates a database directly from each share 4) Create clones of those databases on a nightly basis 5) Run tests directly on those cloned databases
- B. 1) Create a stage in the Production account 2) Create a stage in the QA account that points to the same external object-storage location 3) Create a task that runs nightly to unload each table in the Production account into the stage 4) Use Snowpipe to populate the QA account
- C. 1) Enable replication for each database in the Production account 2) Create replica databases in the QA account 3) Create clones of the replica databases on a nightly basis 4) Run tests directly on those cloned databases
- D. 1) In the Production account, create an external function that connects into the QA account and returns all the data



for one specific table 2) Run the external function as part of a stored procedure that loops through each table in the Production account and populates each table in the QA account

Correct Answer: C

This approach is the least complex because it uses Snowflake's built-in replication feature to copy the data and database objects from the Production account to the QA account. Replication is a fast and efficient way to synchronize data across accounts, regions, and cloud platforms. It also preserves the privileges and metadata of the replicated objects. By creating clones of the replica databases, the QA account can run tests on the cloned data without affecting the original data. Clones are also zero-copy, meaning they do not consume any additional storage space unless the data is modified. This approach does not require any external stages, tasks, Snowpipe, or external functions, which can add complexity and overhead to the data transfer process. References: Introduction to Replication and Failover Replicating Databases Across Multiple Accounts Cloning Considerations

## QUESTION 5

What are purposes for creating a storage integration? (Choose three.)

- A. Control access to Snowflake data using a master encryption key that is maintained in the cloud provider's key management service.
- B. Store a generated identity and access management (IAM) entity for an external cloud provider regardless of the cloud provider that hosts the Snowflake account.
- C. Support multiple external stages using one single Snowflake object.
- D. Avoid supplying credentials when creating a stage or when loading or unloading data.
- E. Create private VPC endpoints that allow direct, secure connectivity between VPCs without traversing the public internet.
- F. Manage credentials from multiple cloud providers in one single Snowflake object.

Correct Answer: BCD

A storage integration is a Snowflake object that stores a generated identity and access management (IAM) entity for an external cloud provider, such as Amazon S3, Google Cloud Storage, or Microsoft Azure Blob Storage. This integration allows Snowflake to read data from and write data to an external storage location referenced in an external stage<sup>1</sup>. One purpose of creating a storage integration is to support multiple external stages using one single Snowflake object. An integration can list buckets (and optional paths) that limit the locations users can specify when creating external stages that use the integration. Note that many external stage objects can reference different buckets and paths and use the same storage integration for authentication<sup>1</sup>. Therefore, option C is correct. Another purpose of creating a storage integration is to avoid supplying credentials when creating a stage or when loading or unloading data. Integrations are named, first-class Snowflake objects that avoid the need for passing explicit cloud provider credentials such as secret keys or access tokens. Integration objects store an IAM user ID, and an administrator in your organization grants the IAM user permissions in the cloud provider account<sup>1</sup>. Therefore, option D is correct. A third purpose of creating a storage integration is to store a generated IAM entity for an external cloud provider regardless of the cloud provider that hosts the Snowflake account. For example, you can create a storage integration for Amazon S3 even if your Snowflake account is hosted on Azure or Google Cloud Platform. This allows you to access data across different cloud platforms using Snowflake<sup>1</sup>. Therefore, option B is correct. Option A is incorrect, because creating a storage integration does not control access to Snowflake data using a master encryption key. Snowflake encrypts all data using a hierarchical key model, and the master encryption key is managed by Snowflake or by the customer using a cloud provider's key management service. This is independent of the storage integration feature<sup>2</sup>. Option E is incorrect, because creating a storage integration does not create private VPC endpoints. Private VPC endpoints are a network configuration option that allow direct, secure connectivity between VPCs without traversing the public internet. This is also independent of



the storage integration feature<sup>3</sup>. Option F is incorrect, because creating a storage integration does not manage credentials from multiple cloud providers in one single Snowflake object. A storage integration is specific to one cloud provider, and you need to create separate integrations for each cloud provider you want to access<sup>4</sup>. References: : Encryption and Decryption : Private Link for Snowflake : CREATE STORAGE INTEGRATION : Option 1: Configuring a Snowflake Storage Integration to Access Amazon S3

[ARA-C01 VCE Dumps](#)

[ARA-C01 Study Guide](#)

[ARA-C01 Braindumps](#)