

Achieve the ARA-C01 Exam Best Results with Help from Snowflake Certified Experts [Q56-Q76]



Achieve the ARA-C01 Exam Best Results with Help from Snowflake Certified Experts Provide ARA-C01 Practice Test Engine for Preparation QUESTION 56

A Snowflake Architect is designing a multi-tenant application strategy for an organization in the Snowflake Data Cloud and is considering using an Account Per Tenant strategy.

Which requirements will be addressed with this approach? (Choose two.)

- * There needs to be fewer objects per tenant.
- * Security and Role-Based Access Control (RBAC) policies must be simple to configure.
- * Compute costs must be optimized.
- * Tenant data shape may be unique per tenant.
- * Storage costs must be optimized.

QUESTION 57

Removing files from a stage after you are done loading the files improves performance when subsequently loading data

- * TRUE
- * FALSE

QUESTION 58

What is a valid object hierarchy when building a Snowflake environment?

- * Account > Database > Schema > Warehouse
- * Organization > Account > Database > Schema > Stage
- * Account > Schema > Table > Stage
- * Organization > Account > Stage > Table > View

QUESTION 59

Which of the below approach results in performance improvement through linear scaling of data ingestion workload?

- * Split large files into recommended range of 10 MB to 100 MB
- * Organize data by granular path
- * Resize virtual warehouse
- * All of the above

QUESTION 60

The endpoint insertFiles of SnowPipe rest API is used to inform Snowflake about the files to be ingested into a table. Select the two statements for this end point.

- * The post can contain at most 5000 files.
- * Each file path given must be <= 1024 bytes long when serialized as UTF-8
- * The post can contain unlimited number of files.

QUESTION 61

You have a table named JSON_TBL which has a variant column JSON_VAR. The json stored in that table looks as below

```
{
  "COURSE_DESC": "SNOWFLAKE CERTIFICATION",
  "COURSE_ID": 1000,
  "DURATION": 2
}
```

if you run a query `SELECT JSON_VAR:Course_id FROM JSON_TBL;` what will it return

- * NULL
- * 1000
- * 2

QUESTION 62

The Data Engineering team at a large manufacturing company needs to engineer data coming from many sources to support a wide

variety of use cases and data consumer requirements which include:

- 1) Finance and Vendor Management team members who require reporting and visualization
 - 2) Data Science team members who require access to raw data for ML model development
 - 3) Sales team members who require engineered and protected data for data monetization
- What Snowflake data modeling approaches will meet these requirements? (Choose two.)
- * Consolidate data in the company's data lake and use EXTERNAL TABLES.
 - * Create a raw database for landing and persisting raw data entering the data pipelines.
 - * Create a set of profile-specific databases that aligns data with usage patterns.
 - * Create a single star schema in a single database to support all consumers' requirements.
 - * Create a Data Vault as the sole data pipeline endpoint and have all consumers directly access the Vault.

QUESTION 63

A user needs access to create materialized view on a schema mydb.myschema.

What is the appropriate command to provide the access?

- * GRANT ROLE MYROLE TO USER USER1; GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO MYROLE;
- * GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER USER1;
- * GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER1;

QUESTION 64

Which feature provides the capability to define an alternate cluster key for a table with an existing cluster key?

- * External table
- * Materialized view
- * Search optimization
- * Result cache

QUESTION 65

Which of the below are securable objects?

- * USER
- * ROLE
- * PRIVILEGE
- * TABLE
- * DATABASE

QUESTION 66

A company is using a Snowflake account in Azure. The account has SAML SSO set up using ADFS as a SCIM identity provider. To validate Private Link connectivity, an Architect performed the following steps:

- * Confirmed Private Link URLs are working by logging in with a username/password account
- * Verified DNS resolution by running nslookups against Private Link URLs

* Validated connectivity using SnowCD

* Disabled public access using a network policy set to use the company's IP address range However, the following error message is received when using SSO to log into the company account:

IP XX.XXX.XX.XX is not allowed to access snowflake. Contact your local security administrator.

What steps should the Architect take to resolve this error and ensure that the account is accessed using only Private Link? (Choose two.)

- * Alter the Azure security integration to use the Private Link URLs.
- * Add the IP address in the error message to the allowed list in the network policy.
- * Generate a new SCIM access token using `system$generate_scim_access_token` and save it to Azure AD.
- * Update the configuration of the Azure AD SSO to use the Private Link URLs.
- * Open a case with Snowflake Support to authorize the Private Link URLs' access to the account.

QUESTION 67

Which statements describe characteristics of the use of materialized views in Snowflake? (Choose two.)

- * They can include ORDER BY clauses.
- * They cannot include nested subqueries.
- * They can include context functions, such as `CURRENT_TIME()`.
- * They can support MIN and MAX aggregates.
- * They can support inner joins, but not outer joins.

QUESTION 68

When a cloned table is replicated to a secondary database, the data also gets replicated in the secondary database

- * TRUE
- * FALSE

QUESTION 69

In the default access control hierarchy, both securityadmin and sysadmin are owned by accountadmin

- * TRUE
- * FALSE

QUESTION 70

Data loading transformation as part of copying data to a table from stage supports selecting data from user stage and named stages(internal and external) only

- * TRUE
- * FALSE

QUESTION 71

To convert JSON null value to SQL null value, you will use

- * `STRIP_NULL_VALUE`
- * `IS_NULL_VALUE`
- * `NULL_IF`

QUESTION 72

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?

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

QUESTION 73

When loading data from stage using COPY INTO, what options can you specify for the ON_ERROR clause?

- * CONTINUE
- * SKIP_FILE
- * ABORT_STATEMENT
- * FAIL

QUESTION 74

SNOWPIPE_AUTO_INGEST is supported for external stages only

- * TRUE
- * FALSE

QUESTION 75

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?

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

QUESTION 76

Which Snowflake data modeling approach is designed for BI queries?

- * 3 NF

- * Star schema
- * Data Vault
- * Snowflake schema

Detailed New ARA-C01 Exam Questions for Concept Clearance: <https://www.validbraindumps.com/ARA-C01-exam-prep.html>