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.leads4pass.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

When loading data from stage using COPY INTO, what options can you specify for the ON_ERROR clause? (Choose three.)

A. CONTINUE

B. SKIP_FILE

C. ABORT_STATEMENT

D. FAIL

Correct Answer: ABC

The ON_ERROR clause is an optional parameter for the COPY INTO command that specifies the behavior of the command when it encounters errors in the files. The ON_ERROR clause can have one of the following values1:

Therefore, options A, B, and C are correct.

References: : COPY INTO

QUESTION 2

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 3

Consider the following COPY command which is loading data with CSV format into a Snowflake table from an internal stage through a data transformation query.

```
copy into home_sales(city, zip, sale_date, price)
from (select t.$1, t.$2, t.$6, t.$7 from @mystage/sales.csv.qz t)
file_format =
(
format_name = mycsvformat
empty_field_as_null = true
field_optionally_enclosed_by = '"'
)
validation_mode = return_all_errors
;
```

This command results in the following error:

SQL compilation error: invalid parameter \\'validation_mode\\'

Assuming the syntax is correct, what is the cause of this error?

A. The VALIDATION_MODE parameter supports COPY statements that load data from external stages only.

B. The VALIDATION_MODE parameter does not support COPY statements with CSV file formats.

C. The VALIDATION_MODE parameter does not support COPY statements that transform data during a load.

D. The value return_all_errors of the option VALIDATION_MODE is causing a compilation error.

Correct Answer: C

The VALIDATION_MODE parameter is used to specify the behavior of the COPY statement when loading data into a table. It is used to specify whether the COPY statement should return an error if any of the rows in the file are invalid or if it should continue loading the valid rows. The VALIDATION_MODE parameter is only supported for COPY statements that load data from external stages1. The query in the question uses a data transformation query to load data from an internal stage. A data transformation query is a query that transforms the data during the load process, such as parsing JSON or XML data, applying functions, or joining with other tables2. According to the documentation, VALIDATION_MODE does not support COPY statements that transform data during a load. If the parameter is specified, the COPY statement returns an error1. Therefore, option C is the correct answer. References: : COPY INTO : Transforming Data During a Load

QUESTION 4

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 stage1. 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 limitthe 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 authentication1. 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 account1. 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 Snowflake1. 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 feature2. 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 feature3. 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 access4. References: : Encryption and Decryption : Private Link for Snowflake : CREATE STORAGE INTEGRATION : Option 1: Configuring a Snowflake Storage Integration to Access Amazon S3

QUESTION 5

A company is storing large numbers of small JSON files (ranging from 1-4 bytes) that are received from IoT devices and sent to a cloud provider. In any given hour, 100,000 files are added to the cloud provider.

What is the MOST cost-effective way to bring this data into a Snowflake table?

A. An external table

B. A pipe

C. A stream

D. A copy command at regular intervals

Correct Answer: B

A pipe is a Snowflake object that continuously loads data from files in a stage (internal or external) into a table. A pipe can be configured to use auto-ingest, which means that Snowflake automatically detects new or modified files in the stage and loads them into the table without any manual intervention1. A pipe is the most cost-effective way to bring large numbers of small JSON files into a Snowflake table, because it minimizes the number of COPY commands executed and the number of micro-partitions created. A pipe can use file aggregation, which means that it can combine multiple small files into a single larger file before loading them into the table. This reduces the load time and the storage

cost of the data2. An external table is a Snowflake object that references data files stored in an external location, such as Amazon S3, Google Cloud Storage, or Microsoft Azure Blob Storage. An external table does not store the data in Snowflake, but only provides a view of the data for querying. An external table is not a cost-effective way to bring data into a Snowflake table, because it does not support file aggregation, and it requires additional network bandwidth and compute resources to

query the external data3.

A stream is a Snowflake object that records the history of changes (inserts, updates, and deletes) made to a table. A stream can be used to consume the changes from a table and apply them to another table or a task. A stream is not a way

to bring data into a Snowflake table, but a way to process the data after it is loaded into a table4.

A copy command is a Snowflake command that loads data from files in a stage into a table. A copy command can be executed manually or scheduled using a task. A copy command is not a cost-effective way to bring large numbers of small

JSON files into a Snowflake table, because it does not support file aggregation, and it may create many micro-partitions that increase the storage cost of the data5.

References: : Pipes : Loading Data Using Snowpipe : External Tables : Streams : COPY INTO

QUESTION 6

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 end 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 end 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 7

What are some of the characteristics of result set caches? (Choose three.)

A. Time Travel queries can be executed against the result set cache.

B. Snowflake persists the data results for 24 hours.

C. Each time persisted results for a query are used, a 24-hour retention period is reset.

D. The data stored in the result cache will contribute to storage costs.

E. The retention period can be reset for a maximum of 31 days.

F. The result set cache is not shared between warehouses.

Correct Answer: BCE

Explanation: Comprehensive and Detailed Explanation: According to the SnowPro Advanced: Architect documents and learning resources, some of the characteristics of result set caches are: Snowflake persists the data results for 24 hours. This means that the result set cache holds the results of every query executed in the past 24 hours, and can be reused if the same query is submitted again and the underlying data has not changed1. Each time persisted results for a query are used, a 24-hour retention period is reset. This means that the result set cache extends the lifetime of the results every time they are reused, up to a maximum of 31 days from the date and time that the query was first executed1. The retention period can be reset for a maximum of 31 days. This means that the result set cache will purge the results after 31 days, regardless of whether they are reused or not. After 31 days, the next time the query is submitted, a new result is generated and persisted1. The other options are incorrect because they are not characteristics of result set caches. Option A is incorrect because Time Travel queries cannot be executed against the result set cache. Time Travel queries use the AS OF clause to access historical data that is stored in the storage layer, not the result set cache2. Option D is incorrect because the data stored in the result set cache does not contribute to storage costs. The result set cache is maintained by the service layer, and does not incur any additional charges1. Option F is incorrect because the result set cache is shared between warehouses. The result set cache is available across virtual warehouses, so query results returned to one user are available to any other user on the system who executes the same query, provided the underlying data has not changed1. References: Using Persisted Query Results | Snowflake Documentation, Time Travel | Snowflake Documentation

---

QUESTION 8

Which organization-related tasks can be performed by the ORGADMIN role? (Choose three.)

A. Changing the name of the organization

B. Creating an account

C. Viewing a list of organization accounts

D. Changing the name of an account

E. Deleting an account

F. Enabling the replication of a database

Correct Answer: BCF

Explanation: According to the SnowPro Advanced: Architect documents and learning resources, the organization-related tasks that can be performed by the ORGADMIN role are: Creating an account in the organization. A user with the ORGADMIN role can use the CREATE ACCOUNT command to create a new account that belongs to the same organization as the current account1. Viewing a list of organization accounts. A user with the ORGADMIN role can use the SHOW ORGANIZATION ACCOUNTS command to view the names and properties of all accounts in the organization2. Alternatively, the user can use the Admin ?Accounts page in the web interface to view the organization name and account names3. Enabling the replication of a database. A user with the ORGADMIN role can use the SYSTEM$GLOBAL_ACCOUNT_SET_PARAMETER function to enable database replication for an account in the organization. This allows the user to replicate databases across accounts in different regions and cloud platforms for data availability and durability4. The other options are incorrect because they are not organization-related tasks that can be performed by the ORGADMIN role. Option A is incorrect because changing the name of the organization is not a task that can be performed by the ORGADMIN role. To change the name of an organization, the user must contact Snowflake Support3. Option D is incorrect because changing the name of an account is not a task that can be performed by the ORGADMIN role. To change the name of an account, the user must contact Snowflake Support5. Option E is incorrect because deleting an account is not a task that can be performed by the ORGADMIN role. To delete an account, the user must contact Snowflake Support. References: CREATE ACCOUNT | Snowflake Documentation, SHOW ORGANIZATION ACCOUNTS | Snowflake Documentation, Getting Started with Organizations | Snowflake Documentation, SYSTEM$GLOBAL_ACCOUNT_SET_PARAMETER | Snowflake Documentation, ALTER ACCOUNT | Snowflake Documentation, [DROP ACCOUNT | Snowflake Documentation]

QUESTION 9

A company has an inbound share set up with eight tables and five secure views. The company plans to make the share part of its production data pipelines.

Which actions can the company take with the inbound share? (Choose two.)

A. Clone a table from a share.

B. Grant modify permissions on the share.

C. Create a table from the shared database.

D. Create additional views inside the shared database.

E. Create a table stream on the shared table.

Correct Answer: AD

Explanation: These two actions are possible with an inbound share, according to the Snowflake documentation and the web search results. An inbound share is a share that is created by another Snowflake account (the provider) and imported into your account (the consumer). An inbound share allows you to access the data shared by the provider, but not to modify or delete it. However, you can perform some actions with the inbound share, such as: Clone a table from a share. You can create a copy of a table from an inbound share using the CREATE TABLE ... CLONE statement. The clone will contain the same data and metadata as the original table, but it will be independent of the share. You can modify or delete the clone as you wish, but it will not reflect any changes made to the original table by the provider1. Create additional views inside the shared database. You can create views on the tables or views from an inbound share using the CREATE VIEW statement. The views will be stored in the shared database, but they will be owned by your account. You can query the views as you would query any other view in your account, but you cannot modify or delete the underlying objects from the share2. The other actions listed are not possible with an inbound share, because they would require modifying the share or the shared objects, which are read-only for the consumer. You cannot grant modify permissions on the share, create a table from the shared database, or create a table stream on the shared table34. References: Cloning Objects from a Share | Snowflake Documentation Creating Views on Shared Data | Snowflake Documentation Importing Data from a Share | Snowflake Documentation Streams on Shared Tables | Snowflake Documentation

---

QUESTION 10

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 usedto 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

---

QUESTION 11

An Architect needs to grant a group of ORDER_ADMIN users the ability to clean old data in an ORDERS table (deleting all records older than 5 years), without granting any privileges on the table. The group\\'s manager (ORDER_MANAGER) has full DELETE privileges on the table.

How can the ORDER_ADMIN role be enabled to perform this data cleanup, without needing the DELETE privilege held by the ORDER_MANAGER role?

A. Create a stored procedure that runs with caller\\'s rights, including the appropriate "> 5 years" business logic, and grant USAGE on this procedure to ORDER_ADMIN. The ORDER_MANAGER role owns the procedure.

B. Create a stored procedure that can be run using both caller\\'s and owner\\'s rights (allowing the user to specify which rights are used during execution), and grant USAGE on this procedure to ORDER_ADMIN. The ORDER_MANAGER role owns the procedure.

C. Create a stored procedure that runs with owner\\'s rights, including the appropriate "> 5 years" business logic, and grant USAGE on this procedure to ORDER_ADMIN. The ORDER_MANAGER role owns the procedure.

D. This scenario would actually not be possible in Snowflake ?any user performing a DELETE on a table requires the DELETE privilege to be granted to the role they are using.

Correct Answer: C

Explanation: This is the correct answer because it allows the ORDER_ADMIN role to perform the data cleanup without needing the DELETE privilege on the ORDERS table. A stored procedure is a feature that allows scheduling and executing SQL statements or stored procedures in Snowflake. A stored procedure can run with either the caller\\'s rights or the owner\\'s rights. A caller\\'s rights stored procedure runs with the privileges of the role that called the stored procedure, while an owner\\'s rights stored procedure runs with the privileges of the role that created the stored procedure. By creating a stored procedure that runs with owner\\'s rights, the ORDER_MANAGER role can delegate the specific task of deleting old data to the ORDER_ADMIN role, without granting the ORDER_ADMIN role more general privileges on the ORDERS table. The stored procedure must include the appropriate business logic to delete only the records older than 5 years, and the ORDER_MANAGER role must grant the USAGE privilege on the stored procedure to the ORDER_ADMIN role. The ORDER_ADMIN role can then execute the stored procedure to perform the data cleanup12. References: Snowflake Documentation: Stored Procedures Snowflake Documentation: Understanding Caller\\'s Rights and Owner\\'s Rights Stored Procedures

---

QUESTION 12

What built-in Snowflake features make use of the change tracking metadata for a table? (Choose two.)

A. The MERGE command

B. The UPSERT command

C. The CHANGES clause

D. A STREAM object

E. TheCHANGE_DATA_CAPTURE command

Correct Answer: CD

Explanation: The built-in Snowflake features that make use of the change tracking metadata for a table are the CHANGES clause and a STREAM object. The CHANGES clause enables querying the change tracking metadata for a table or view within a specified interval of time without having to create a stream with an explicit transactional offset1. A STREAM object records data manipulation language (DML) changes made to tables, including inserts, updates, and deletes, as well as metadata about each change, so that actions can be taken using the changed data. This process is referred to as change data capture (CDC)2. The other options are incorrect because they do not make use of the change tracking metadata for a table. The MERGE command performs insert, update, or delete operations on a target table based on the results of a join with a source table3. The UPSERT command is not a valid Snowflake command. The CHANGE_DATA_CAPTURE command is not a valid Snowflake command. References: CHANGES | Snowflake Documentation, Change Tracking Using Table Streams | Snowflake Documentation, MERGE | Snowflake Documentation

QUESTION 13

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

QUESTION 14

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 15

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.)

A. Consolidate data in the company\\'s data lake and use EXTERNAL TABLES.

B. Create a raw database for landing and persisting raw data entering the data pipelines.

C. Create a set of profile-specific databases that aligns data with usage patterns.

D. Create a single star schema in a single database to support all consumers\\' requirements.

E. Create a Data Vault as the sole data pipeline endpoint and have all consumers directly access the Vault.

Correct Answer: BC

Explanation: These two approaches are recommended by Snowflake for data modeling in a data lake scenario. Creating a raw database allows the data engineering team to ingest data from various sources without any transformation or cleansing, preserving the original data quality and format. This enables the data science team to access the raw data for ML model development. Creating a set of profile-specific databases allows the data engineering team to apply different transformations and optimizations for different use cases and data consumer requirements. For example, the finance and vendor management team can access a dimensional database that supports reporting and visualization, while the sales team can access a secure database that supports data monetization. References: Snowflake Data Lake Architecture | Snowflake Documentation Snowflake Data Lake Best Practices | Snowflake Documentation