# 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.pass2lead.com/ara-c01.html**

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Snowflake Official Exam Center

**QUESTION 1**

A media company needs a data pipeline that will ingest customer review data into a Snowflake table, and apply some transformations. The company also needs to use Amazon Comprehend to do sentiment analysis and make the deidentified final data set available publicly for advertising companies who use different cloud providers in different regions.

The data pipeline needs to run continuously ang efficiently as new records arrive in the object storage leveraging event notifications. Also, the operational complexity, maintenance of the infrastructure, including platform upgrades and security, and the development effort should be minimal.

Which design will meet these requirements?

A. Ingest the data using COPY INTO and use streams and tasks to orchestrate transformations. Export the data into Amazon S3 to do model inference with Amazon Comprehend and ingest the data back into a Snowflake table. Then create a listing in the Snowflake Marketplace to make the data available to other companies.

B. Ingest the data using Snowpipe and use streams and tasks to orchestrate transformations. Create an external function to do model inference with Amazon Comprehend and write the final records to a Snowflake table. Then create a listing in the Snowflake Marketplace to make the data available to other companies.

C. Ingest the data into Snowflake using Amazon EMR and PySpark using the Snowflake Spark connector. Apply transformations using another Spark job. Develop a python program to do model inference by leveraging the Amazon Comprehend text analysis API. Then write the results to a Snowflake table and create a listing in the Snowflake Marketplace to make the data available to other companies.

D. Ingest the data using Snowpipe and use streams and tasks to orchestrate transformations. Export the data into Amazon S3 to do model inference with Amazon Comprehend and ingest the data back into a Snowflake table. Then create a listing in the Snowflake Marketplace to make the data available to other companies.

Correct Answer: B

Explanation: This design meets all the requirements for the data pipeline. Snowpipe is a feature that enables continuous data loading into Snowflake from object storage using event notifications. It is efficient, scalable, and serverless,

meaning it does not require any infrastructure or maintenance from the user. Streams and tasks are features that enable automated data pipelines within Snowflake, using change data capture and scheduled execution. They are also efficient,

scalable, and serverless, and they simplify the data transformation process. External functions are functions that can invoke external services or APIs from within Snowflake. They can be used to integrate with Amazon Comprehend and

perform sentiment analysis on the data. The results can be written back to a Snowflake table using standard SQL commands. Snowflake Marketplace is a platform that allows data providers to share data with data consumers across different

accounts, regions, and cloud platforms. It is a secure and easy way to make data publicly available to other companies.

References:

Snowpipe Overview | Snowflake Documentation

Introduction to Data Pipelines | Snowflake Documentation External Functions Overview | Snowflake Documentation
Snowflake Data Marketplace Overview | Snowflake Documentation

**QUESTION 2**

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.)

A. Alter the Azure security integration to use the Private Link URLs.

B. Add the IP address in the error message to the allowed list in the network policy.

C. Generate a new SCIM access token using system$generate_scim_access_token and save it to Azure AD.

D. Update the configuration of the Azure AD SSO to use the Private Link URLs.

E. Open a case with Snowflake Support to authorize the Private Link URLs\\' access to the account.

Correct Answer: BD

Explanation: The error message indicates that the IP address in the error message is not allowed to access Snowflake because it is not in the allowed list of the network policy. The network policy is a feature that allows restricting access to Snowflake based on IP addresses or ranges. To resolve this error, the Architect should take the following steps: Add the IP address in the error message to the allowed list in the network policy. This will allow the IP address to access Snowflake using the Private Link URLs. Alternatively, the Architect can disable the network policy if it is not required for security reasons. Update the configuration of the Azure AD SSO to use the Private Link URLs. This will ensure that the SSO authentication process uses the Private Link URLs instead of the public URLs. The configuration can be updated by following the steps in the Azure documentation1. These two steps should resolve the error and ensure that the account is accessed using only Private Link. The other options are not necessary or relevant for this scenario. Altering the Azure security integration to use the Private Link URLs is not required because the security integration is used for SCIM provisioning, not for SSO authentication. Generating a new SCIM access token using system$generate_scim_access_token and saving it to Azure AD is not required because the SCIM access token is used for SCIM provisioning, not for SSO authentication. Opening a case with Snowflake Support to authorize the Private Link URLs\\' access to the account is not required because the authorization can be done by the account administrator using the SYSTEM$AUTHORIZE_PRIVATELINK function2.

**QUESTION 3**

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 4**

What is a characteristic of loading data into Snowflake using the Snowflake Connector for Kafka?

A. The Connector only works in Snowflake regions that use AWS infrastructure.

B. The Connector works with all file formats, including text, JSON, Avro, Ore, Parquet, and XML.

C. The Connector creates and manages its own stage, file format, and pipe objects.

D. Loads using the Connector will have lower latency than Snowpipe and will ingest data in real time.

Correct Answer: C

Explanation: According to the SnowPro Advanced: Architect documents and learning resources, a characteristic of loading data into Snowflake using the Snowflake Connector for Kafka is that the Connector creates and manages its own stage, file format, and pipe objects. The stage is an internal stage that is used to store the data files from the Kafka topics. The file format is a JSON or Avro file format that is used to parse the data files. The pipe is a Snowpipe object that is used to load the data files into the Snowflake table. The Connector automatically creates and configures these objects based on the Kafka configuration properties, and handles the cleanup and maintenance of these objects1. The other options are incorrect because they are not characteristics of loading data into Snowflake using the Snowflake Connector for Kafka. Option A is incorrect because the Connector works in Snowflake regions that use any cloud infrastructure, not just AWS. The Connector supports AWS, Azure, and Google Cloud platforms, and can load data across different regions and cloud platforms using data replication2. Option B is incorrect because the Connector does not work with all file formats, only JSON and Avro. The Connector expects the data in the Kafka topics to be in JSON or Avro format, and parses the data accordingly. Other file formats, such as text, ORC, Parquet, or XML, are not supported by the Connector3. Option D is incorrect because loads using the Connector do not have lower latency than Snowpipe, and do not ingest data in real time. The Connector uses Snowpipe to load data into Snowflake, and inherits the same latency and performance characteristics of Snowpipe. The Connector does not provide real-time ingestion, but near real-time ingestion, depending on the frequency and size of the data files4. References: Installing and Configuring the Kafka Connector | Snowflake Documentation, Sharing Data Across Regions and Cloud Platforms | Snowflake Documentation, Overview of the Kafka Connector | Snowflake Documentation, Using Snowflake Connector for Kafka With Snowpipe Streaming | Snowflake Documentation

**QUESTION 5**

An Architect has been asked to clone schema STAGING as it looked one week ago, Tuesday June 1st at 8:00 AM, to recover some objects.

The STAGING schema has 50 days of retention.

The Architect runs the following statement:

CREATE SCHEMA STAGING_CLONE CLONE STAGING at (timestamp => \\'2021-06-01 08:00:00\\');

The Architect receives the following error: Time travel data is not available for schema STAGING. The requested time is either beyond the allowed time travel period or before the object creation time.

The Architect then checks the schema history and sees the following:

CREATED_ON|NAME|DROPPED_ON

2021-06-02 23:00:00 | STAGING | NULL

2021-05-01 10:00:00 | STAGING | 2021-06-02 23:00:00

How can cloning the STAGING schema be achieved?

A. Undrop the STAGING schema and then rerun the CLONE statement.

B. Modify the statement: CREATE SCHEMA STAGING_CLONE CLONE STAGING at (timestamp => \\'2021-05-01

10:00:00\\');

C. Rename the STAGING schema and perform an UNDROP to retrieve the previous STAGING schema version, then run the CLONE statement.

D. Cloning cannot be accomplished because the STAGING schema version was not active during the proposed Time Travel time period.

Correct Answer: C

The error message indicates that the schema STAGING does not have time travel data available for the requested timestamp, because the current version of the schema was created on2021-06-02 23:00:00, which is after the timestamp of 2021-06-01 08:00:00. Therefore, the CLONE statement cannot access the historical data of the schema at that point in time. Option A is incorrect, because undropping the STAGING schema will not restore the previous version of the schema that was active on 2021-06-01 08:00:00. Instead, it will create a new version of the schema with the same name and no data or objects. Option B is incorrect, because modifying the timestamp to 2021-05-01 10:00:00 will not clone the schema as it looked one week ago, but as it looked when it was first created. This may not reflect the desired state of the schema and its objects. Option C is correct, because renaming the STAGING schema and performing an UNDROP to retrieve the previous STAGING schema version will restore the schema that was dropped on 2021-06-02 23:00:00. This schema has time travel data available for the requested timestamp of 2021-06-01 08:00:00, and can be cloned using the CLONE statement. Option D is incorrect, because cloning can be accomplished by using the UNDROP command to access the previous version of the schema that was active during the proposed time travel period. References: : Cloning Considerations : Understanding and Using Time Travel : CREATE ... CLONE

ARA-C01 Practice Test ARA-C01 Study Guide ARA-C01 Braindumps