

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

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



QUESTION 1

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 2

Which security, governance, and data protection features require, at a MINIMUM, the Business Critical edition of Snowflake? (Choose two.)

- A. Extended Time Travel (up to 90 days)
- B. Customer-managed encryption keys through Tri-Secret Secure

- C. Periodic rekeying of encrypted data
- D. AWS, Azure, or Google Cloud private connectivity to Snowflake
- E. Federated authentication and SSO

Correct Answer: BD

Explanation: According to the SnowPro Advanced: Architect documents and learning resources, the security, governance, and data protection features that require, at a minimum, the Business Critical edition of Snowflake are: Customer-managed encryption keys through Tri-Secret Secure. This feature allows customers to manage their own encryption keys for data at rest in Snowflake, using a combination of three secrets: a master key, a service key, and a security password. This provides an additional layer of security and control over the data encryption and decryption process¹. Periodic rekeying of encrypted data. This feature allows customers to periodically rotate the encryption keys for data at rest in Snowflake, using either Snowflake- managed keys or customer-managed keys. This enhances the security and protection of the data by reducing the risk of key compromise or exposure². The other options are incorrect because they do not require the Business Critical edition of Snowflake. Option A is incorrect because extended Time Travel (up to 90 days) is available with the Enterprise edition of Snowflake³. Option D is incorrect because AWS, Azure, or Google Cloud private connectivity to Snowflake is available with the Standard edition of Snowflake⁴. Option E is incorrect because federated authentication and SSO are available with the Standard edition of Snowflake⁵. References: Tri-Secret Secure | Snowflake Documentation, Periodic Rekeying of Encrypted Data | Snowflake Documentation, Snowflake Editions | Snowflake Documentation, Snowflake Network Policies | Snowflake Documentation, Configuring Federated Authentication and SSO | Snowflake Documentation

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

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

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

Correct Answer: BD

Explanation: According to the Snowflake documentation, materialized views have some limitations on the query specification that defines them. One of these limitations is that they cannot include nested subqueries, such as subqueries in the `FROM` clause or scalar subqueries in the `SELECT` list. Another limitation is that they cannot include `ORDER BY` clauses, context functions (such as `CURRENT_TIME()`), or outer joins. However, materialized views can support `MIN` and `MAX` aggregates, as well as other aggregate functions, such as `SUM`, `COUNT`, and `AVG`. References: Limitations on Creating Materialized Views | Snowflake Documentation Working with Materialized Views | Snowflake Documentation

QUESTION 5

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

[ARA-C01 VCE Dumps](#)

[ARA-C01 Practice Test](#)

[ARA-C01 Study Guide](#)