Snowflake ARA-C01 Mock Exam

Shared by Walters on 17-06-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Which columns can be included in an external table schema? (Select THREE).

Options:

A- VALUE
B- METADATASROW_ID
C- METADATASISUPDATE
D- METADAT A$ FILENAME
E- METADATAS FILE_ROW_NUMBER
F- METADATASEXTERNAL TABLE PARTITION

Answer:

A, D, E

Explanation:

An external table schema defines the columns and data types of the data stored in an external stage. All external tables include the following columns by default:

VALUE: A VARIANT type column that represents a single row in the external file.

METADATA$FILENAME: A pseudocolumn that identifies the name of each staged data file included in the external table, including its path in the stage.

METADATA$FILE_ROW_NUMBER: A pseudocolumn that shows the row number for each record in a staged data file.

You can also create additional virtual columns as expressions using the VALUE column and/or the pseudocolumns. However, the following columns are not valid for external tables and cannot be included in the schema:

METADATASROW_ID: This column is only available for internal tables and shows the unique identifier for each row in the table.

METADATASISUPDATE: This column is only available for internal tables and shows whether the row was inserted or updated by a merge operation.

METADATASEXTERNAL TABLE PARTITION: This column is not a valid column name and does not exist in Snowflake.

Question 2

Question Type: MultipleChoice

What is the MOST efficient way to design an environment where data retention is not considered critical, and customization needs are to be kept to a minimum?

Options:

A- Use a transient database.
B- Use a transient schema.
C- Use a transient table.
D- Use a temporary table.

Answer:

A

Explanation:

Transient databases in Snowflake are designed for situations where data retention is not critical, and they do not have the fail-safe period that regular databases have. This means that data in a transient database is not recoverable after the Time Travel retention period. Using a transient database is efficient because it minimizes storage costs while still providing most functionalities of a standard database without the overhead of data protection features that are not needed when data retention is not a concern.

Question 3

Question Type: MultipleChoice

The diagram shows the process flow for Snowpipe auto-ingest with Amazon Simple Notification Service (SNS) with the following steps:

Step 1: Data files are loaded in a stage.

Step 2: An Amazon S3 event notification, published by SNS, informs Snowpipe --- by way of Amazon Simple Queue Service (SQS) - that files are ready to load. Snowpipe copies the files into a queue.

Step 3: A Snowflake-provided virtual warehouse loads data from the queued files into the target table based on parameters defined in the specified pipe.

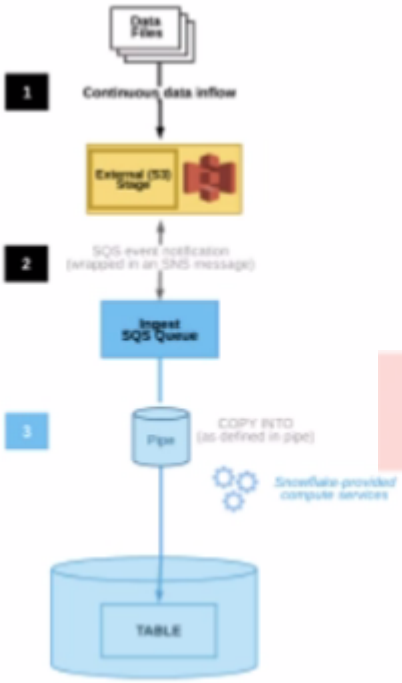

If an AWS Administrator accidentally deletes the SQS subscription to the SNS topic in Step 2, what will happen to the pipe that references the topic to receive event messages from Amazon S3?

Options:

A- The pipe will continue to receive the messages as Snowflake will automatically restore the subscription to the same SNS topic and will recreate the pipe by specifying the same SNS topic name in the pipe definition.

B- The pipe will no longer be able to receive the messages and the user must wait for 24 hours from the time when the SNS topic subscription was deleted. Pipe recreation is not required as the pipe will reuse the same subscription to the existing SNS topic after 24 hours.

C- The pipe will continue to receive the messages as Snowflake will automatically restore the subscription by creating a new SNS topic. Snowflake will then recreate the pipe by specifying the new SNS topic name in the pipe definition.

D- The pipe will no longer be able to receive the messages. To restore the system immediately, the user needs to manually create a new SNS topic with a different name and then recreate the pipe by specifying the new SNS topic name in the pipe definition.

Answer:

D

Explanation:

If an AWS Administrator accidentally deletes the SQS subscription to the SNS topic in Step 2, the pipe that references the topic to receive event messages from Amazon S3 will no longer be able to receive the messages. This is because the SQS subscription is the link between the SNS topic and the Snowpipe notification channel. Without the subscription, the SNS topic will not be able to send notifications to the Snowpipe queue, and the pipe will not be triggered to load the new files. To restore the system immediately, the user needs to manually create a new SNS topic with a different name and then recreate the pipe by specifying the new SNS topic name in the pipe definition. This will create a new notification channel and a new SQS subscription for the pipe. Alternatively, the user can also recreate the SQS subscription to the existing SNS topic and then alter the pipe to use the same SNS topic name in the pipe definition. This will also restore the notification channel and the pipe functionality.Reference:

Automating Snowpipe for Amazon S3

Enabling Snowpipe Error Notifications for Amazon SNS

HowTo: Configuration steps for Snowpipe Auto-Ingest with AWS S3 Stages

Question 4

Question Type: MultipleChoice

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?

Options:

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.

Answer:

C

Explanation:

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 us-east-1 region with the PARTNERB account in Azure East US 2 region, the following steps must be performed:

Create a new account (called AZABC123) in Azure East US 2 region. This account will act as a bridge between the source and the target accounts.The new account must be linked to the ACCOUNTA account using an organization2.

From the ACCOUNTA account, replicate the MARKET_DB database to the AZABC123 account using the account replication feature.This will create a secondary database in the AZABC123 account that is a replica of the primary database in the ACCOUNTA account3.

From the AZABC123 account, set up the data sharing to the PARTNERB account using the share replication feature. This will create a share of the secondary database in the AZABC123 account and grant access to the PARTNERB account.The PARTNERB account can then create a database from the share and query the data4.

Therefore, option C is the correct answer.

Question 5

Question Type: MultipleChoice

Which system functions does Snowflake provide to monitor clustering information within a table (Choose two.)

Options:

A- SYSTEM$CLUSTERING_INFORMATION
B- SYSTEM$CLUSTERING_USAGE
C- SYSTEM$CLUSTERING_DEPTH
D- SYSTEM$CLUSTERING_KEYS
E- SYSTEM$CLUSTERING_PERCENT

Answer:

A, C

Explanation:

According to the Snowflake documentation, these two system functions are provided by Snowflake to monitor clustering information within a table. A system function is a type of function that allows executing actions or returning information about the system. A clustering key is a feature that allows organizing data across micro-partitions based on one or more columns in the table. Clustering can improve query performance by reducing the number of files to scan.

SYSTEM$CLUSTERING_INFORMATION is a system function that returns clustering information, including average clustering depth, for a table based on one or more columns in the table. The function takes a table name and an optional column name or expression as arguments, and returns a JSON string with the clustering information.The clustering information includes the cluster by keys, the total partition count, the total constant partition count, the average overlaps, and the average depth1.

SYSTEM$CLUSTERING_DEPTH is a system function that returns the clustering depth for a table based on one or more columns in the table. The function takes a table name and an optional column name or expression as arguments, and returns an integer value with the clustering depth. The clustering depth is the maximum number of overlapping micro-partitions for any micro-partition in the table.A lower clustering depth indicates a better clustering2.

SYSTEM$CLUSTERING_INFORMATION | Snowflake Documentation

SYSTEM$CLUSTERING_DEPTH | Snowflake Documentation

Question 6

Question Type: MultipleChoice

What transformations are supported in the below SQL statement? (Select THREE).

CREATE PIPE ... AS COPY ... FROM (...)

Options:

A- Data can be filtered by an optional where clause.
B- Columns can be reordered.
C- Columns can be omitted.
D- Type casts are supported.
E- Incoming data can be joined with other tables.
F- The ON ERROR - ABORT statement command can be used.

Answer:

A, B, C

Explanation:

The SQL statement is a command for creating a pipe in Snowflake, which is an object that defines the COPY INTO <table> statement used by Snowpipe to load data from an ingestion queue into tables1.The statement uses a subquery in the FROM clause to transform the data from the staged files before loading it into the table2.

The transformations supported in the subquery are as follows2:

Data can be filtered by an optional WHERE clause, which specifies a condition that must be satisfied by the rows returned by the subquery. For example:

SQLAI-generated code. Review and use carefully.More info on FAQ.

```
create pipe mypipe as
copy into mytable
from (
select * from @mystage
where col1 = 'A' and col2 > 10
);
```

Columns can be reordered, which means changing the order of the columns in the subquery to match the order of the columns in the target table. For example:

SQLAI-generated code. Review and use carefully.More info on FAQ.

```
create pipe mypipe as
copy into mytable (col1, col2, col3)
```

```
from (
select col3, col1, col2 from @mystage
);
```

Columns can be omitted, which means excluding some columns from the subquery that are not needed in the target table. For example:

SQLAI-generated code. Review and use carefully.More info on FAQ.

```
create pipe mypipe as
copy into mytable (col1, col2)
from (
select col1, col2 from @mystage
);
```

The other options are not supported in the subquery because2:

Type casts are not supported, which means changing the data type of a column in the subquery. For example, the following statement will cause an error:

SQLAI-generated code. Review and use carefully.More info on FAQ.

```
create pipe mypipe as
copy into mytable (col1, col2)
from (
select col1::date, col2 from @mystage
);
```

Incoming data can not be joined with other tables, which means combining the data from the staged files with the data from another table in the subquery. For example, the following statement will cause an error:

SQLAI-generated code. Review and use carefully.More info on FAQ.

```
create pipe mypipe as
copy into mytable (col1, col2, col3)
from (
select s.col1, s.col2, t.col3 from @mystage s
join othertable t on s.col1 = t.col1
```

);

The ON ERROR - ABORT statement command can not be used, which means aborting the entire load operation if any error occurs. This command can only be used in the COPY INTO <table> statement, not in the subquery. For example, the following statement will cause an error:

SQLAI-generated code. Review and use carefully.More info on FAQ.

```
create pipe mypipe as

copy into mytable

from (

select * from @mystage

on error abort

);
```

1: CREATE PIPE | Snowflake Documentation

2: Transforming Data During a Load | Snowflake Documentation

P2P exams