

# ARA-C01考試資訊 & ARA-C01題庫更新資訊



順便提一下，可以從雲存儲中下載NewDumps ARA-C01考試題庫的完整版：<https://drive.google.com/open?id=1tyBBXeY62GhkpgLWkwj9p6XLwo4XXEiL>

NewDumps的專家團隊利用他們的經驗和知識終於研究出了關於Snowflake ARA-C01 認證考試的培訓資料。我們的Snowflake ARA-C01 認證考試培訓資料很受客戶歡迎，這是NewDumps的專家團隊勤勞勞動的結果。他們研究出來的模擬測試題及答案有很高的品質，和真實的考試題目有95%的相似性，是很值得你依賴的。如果你使用了NewDumps的培訓工具，你可以100%通過你的第一次參加的Snowflake ARA-C01認證考試。

NewDumps剛剛發布了最新的ARA-C01認證考試所有更新的問題及答案，來確保您考試成功通過。我們提供最新的PDF和軟件版本的問題和答案，可以保證考生的ARA-C01考試100%通過。在我們的網站上，您將獲得我們提供的Snowflake ARA-C01免費的PDF版本的DEMO試用，您會發現這絕對是最值得信賴的學習資料。對於擁有高命中率的Snowflake ARA-C01考古題，還在等什么，趕快下載最新的題庫資料來準備考試吧！

>> ARA-C01考試資訊 <<

## ARA-C01題庫更新資訊 - ARA-C01考題寶典

我們都是平平凡凡的普通人，有時候所學的所掌握的東西沒有那麼容易徹底的吸收，所以經常忘記，當我們需要時就拼命的補習，當你看到NewDumps Snowflake的ARA-C01考試培訓資料是，你才明白這是你必須要購買的，它可以讓你毫不費力的通過考試，也可以讓你不那么努力的補習，相信NewDumps，相信它讓你看到你的未來美好的樣子，再苦再難，只要NewDumps還在，總會找到希望的光明。

## 最新的 SnowPro Advanced Certification ARA-C01 免費考試真題 (Q148-Q153):

### 問題 #148

A company is designing a process for importing a large amount of IoT JSON data from cloud storage into Snowflake. New sets of IoT data get generated and uploaded approximately every 5 minutes.

Once the IoT data is in Snowflake, the company needs up-to-date information from an external vendor to join to the data. This data is then presented to users through a dashboard that shows different levels of aggregation.

The external vendor is a Snowflake customer.

What solution will MINIMIZE complexity and MAXIMIZE performance?

- A. 1. Create a Snowpipe to bring the JSON data into Snowflake.  
2. Use streams and tasks to trigger a transformation procedure when new JSON data arrives.  
3. Ask the vendor to expose an API so an external function call can be made to join the vendor's data back to the IoT data in a transformation procedure.  
4. Create materialized views over the larger dataset to perform the aggregations required by the dashboard.  
5. Give the materialized views access to the dashboard tool.
- B. 1. Create an external table over the JSON data in cloud storage.  
2. Create a task that runs every 5 minutes to run a transformation procedure on new data, based on a saved timestamp.  
3. Ask the vendor to expose an API so an external function can be used to generate a call to join the data back to the IoT data in the transformation procedure.  
4. Give the transformed table access to the dashboard tool.  
5. Perform the aggregations on the dashboard tool.
- C. 1. Create an external table over the JSON data in cloud storage.  
2. Create a task that runs every 5 minutes to run a transformation procedure on new data based on a saved timestamp.  
3. Ask the vendor to create a data share with the required data that can be imported into the company's Snowflake account.  
4. Join the vendor's data back to the IoT data using a transformation procedure.  
5. Create views over the larger dataset to perform the aggregations required by the dashboard.  
6. Give the views access to the dashboard tool.
- D. 1. Create a Snowpipe to bring the JSON data into Snowflake.  
2. Use streams and tasks to trigger a transformation procedure when new JSON data arrives.  
3. Ask the vendor to create a data share with the required data that is then imported into the Snowflake account.  
4. Join the vendor's data back to the IoT data in a transformation procedure  
5. Create materialized views over the larger dataset to perform the aggregations required by the dashboard.  
6. Give the materialized views access to the dashboard tool.

答案： D

解題說明：

Using Snowpipe for continuous, automated data ingestion minimizes the need for manual intervention and ensures that data is available in Snowflake promptly after it is generated. Leveraging Snowflake's data sharing capabilities allows for efficient and secure access to the vendor's data without the need for complex API integrations. Materialized views provide pre-aggregated data for fast access, which is ideal for dashboards that require high performance<sup>1234</sup>.

References =

\*Snowflake Documentation on Snowpipe<sup>4</sup>

\*Snowflake Documentation on Secure Data Sharing<sup>2</sup>

\*Best Practices for Data Ingestion with Snowflake<sup>1</sup>

### 問題 #149

You are creating a TASK to query a table streams created on the raw table and insert subsets of rows into multiple tables. You are following the below steps, but when you reached the step to resume the task, you received an error message as below.

Why is this error thrown and who can give you the required privilege?

```
51 -- Resume both tasks.  
52 alter task raw_to_names resume;  
53 alter task raw_to_visits resume;
```

results Data Preview

Query ID	SQL	38ms
----------	-----	------

Cannot execute task , EXECUTE TASK privilege must be granted to owner role

Steps to be followed to get this error

- Create a landing table to store raw JSON data.
- Snowpipe could load data into this table. create or replace table raw (var variant);
- Create a stream to capture inserts to the landing table.
- A task will consume a set of columns from this stream. create or replace stream rawstream1 on table raw;
- Create a second stream to capture inserts to the landing table.
- A second task will consume another set of columns from this stream. create or replace stream rawstream2 on table raw;
- Create a table that stores the names of office visitors identified in the raw data. create or replace table names (id int, first\_name string, last\_name string);

```

-- Create a table that stores the visitation dates of office visitors identified in the raw data.
create or replace table visits (id int, dt date);
-- Create a task that inserts new name records from the rawstream1 stream into the names table
-- every minute when the stream contains records.
-- Replace the 'mywh' warehouse with a warehouse that your role has USAGE privilege on. create or replace task raw_to_names
warehouse = etl_wh
schedule = '1 minute'
when
system$stream_has_data('rawstream1')
as
merge into names n
using (select var:id id, var:fname fname, var:lname lname from rawstream1) r1 on n.id = to_number(r1.id)
when matched then update set n.first_name = r1.fname, n.last_name = r1.lname
when not matched then insert (id, first_name, last_name) values (r1.id, r1.fname, r1.lname)
;
-- Create another task that merges visitation records from the rawstream1 stream into the visits table
-- every minute when the stream contains records.
-- Records with new IDs are inserted into the visits table;
-- Records with IDs that exist in the visits table update the DT column in the table.
-- Replace the 'mywh' warehouse with a warehouse that your role has USAGE privilege on. create or replace task raw_to_visits
warehouse = etl_wh schedule = '1 minute' when
system$stream_has_data('rawstream2') as
merge into visits v
using (select var:id id, var:visit_dt visit_dt from rawstream2) r2 on v.id = to_number(r2.id) when matched then update set v.dt =
r2.visit_dt
when not matched then insert (id, dt) values (r2.id, r2.visit_dt);
-- Resume both tasks.
alter task raw_to_names resume;

```

- A. The role used to resume the task does not have EXECUTE TASK privilege. Only TASK OWNER can provide that privilege to the role.
- B. The role used to resume the task does not have EXECUTE TASK privilege. Both SECURITYADMIN and ACCOUNTADMIN can provide that privilege to the role.
- C. The role used to resume the task does not have EXECUTE TASK privilege. Only ACCOUNTADMIN can provide that privilege to the role.

答案： C

### 問題 #150

When loading data into a table that captures the load time in a column with a default value of either CURRENT\_TIME () or CURRENT\_TIMESTAMP() what will occur?

- A. All rows loaded using a specific COPY statement will have the same timestamp value.
- B. Any rows loaded using a specific COPY statement will have varying timestamps based on when the rows were read from the source.
- C. Any rows loaded using a specific COPY statement will have varying timestamps based on when the rows were created in the source.
- D. All rows loaded using a specific COPY statement will have varying timestamps based on when the rows were inserted.

答案： A

解題說明：

Explanation

According to the Snowflake documentation, when loading data into a table that captures the load time in a column with a default value of either CURRENT\_TIME () or CURRENT\_TIMESTAMP(), the default value is evaluated once per COPY statement, not once per row. Therefore, all rows loaded using a specific COPY statement will have the same timestamp value. This behavior ensures that the timestamp value reflects the time when the data was loaded into the table, not when the data was read from the source or created in the source.

References:

\* Snowflake Documentation: Loading Data into Tables with Default Values

\* Snowflake Documentation: COPY INTO table

### 問題 #151

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. 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.
- 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. 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.

答案： C

解題說明：

\* 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 organization<sup>1</sup>.

\* 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 organization<sup>2</sup>.

\* 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 account<sup>3</sup>.

\* 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 data<sup>4</sup>.

\* Therefore, option C is the correct answer.

Replicating Shares Across Regions and Cloud Platforms : Working with Organizations and Accounts : Replicating Databases Across Multiple Accounts : Replicating Shares Across Multiple Accounts

### 問題 #152

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 should use a storage integration for the external stage.
- 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 must have a Snowflake account in each cloud region to be able to ingest data to that account.

答案： A

解題說明：

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 Snowflake<sup>12</sup>.

Reference:



id=1tyBBXeY62GhkpgLWkwj9p6XLwo4XXEiL