

DUMPS ARENA

SnowPro Advanced: Architect Certification Exam

Snowflake ARA-C01

Version Demo

Total Demo Questions: 10

Total Premium Questions: 65

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

A company's daily Snowflake workload consists of a huge number of concurrent queries triggered between 9pm and 11pm. At the individual level, these queries are smaller statements that get completed within a short time period.

What configuration can the company's Architect implement to enhance the performance of this workload? (Choose two.)

- A. Enable a multi-clustered virtual warehouse in maximized mode during the workload duration.
- B. Set the MAX_CONCURRENCY_LEVEL to a higher value than its default value of 8 at the virtual warehouse level.
- C. Increase the size of the virtual warehouse to size X-Large.
- D. Reduce the amount of data that is being processed through this workload.
- E. Set the connection timeout to a higher value than its default.

ANSWER: A C**QUESTION NO: 2**

What is a valid object hierarchy when building a Snowflake environment?

- A. Account --> Database --> Schema --> Warehouse
- B. Organization --> Account --> Database --> Schema --> Stage
- C. Account --> Schema > Table --> Stage
- D. Organization --> Account --> Stage --> Table --> View

ANSWER: B**QUESTION NO: 3**

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.

ANSWER: B C E**QUESTION NO: 4**

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.

ANSWER: C**QUESTION NO: 5**

Which of the following are characteristics of how row access policies can be applied to external tables? (Choose three.)

- A.** An external table can be created with a row access policy, and the policy can be applied to the VALUE column.
- B.** A row access policy can be applied to the VALUE column of an existing external table.
- C.** A row access policy cannot be directly added to a virtual column of an external table.
- D.** External tables are supported as mapping tables in a row access policy.
- E.** While cloning a database, both the row access policy and the external table will be cloned.
- F.** A row access policy cannot be applied to a view created on top of an external table.

ANSWER: A B C**QUESTION NO: 6**

What integration object should be used to place restrictions on where data may be exported?

- A. Stage integration
- B. Security integration
- C. Storage integration
- D. API integration

ANSWER: C

QUESTION NO: 7

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 de-identified final data set available publicly for advertising companies who use different cloud providers in different regions.

The data pipeline needs to run continuously and 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.

ANSWER: B

QUESTION NO: 8

An Architect entered the following commands in sequence:

```
CREATE DATABASE SANDBOX;  
CREATE ROLE INTERN;  
CREATE TABLE SANDBOX.PUBLIC.AGENDA (ID INT, ITEMS STRING);  
GRANT SELECT ON ALL TABLES IN SCHEMA SANDBOX.PUBLIC TO ROLE INTERN;  
GRANT ROLE INTERN TO USER USER1;
```

USER1 cannot find the table.

Which of the following commands does the Architect need to run for USER1 to find the tables using the Principle of Least Privilege? (Choose two.)

- A. GRANT ROLE PUBLIC TO ROLE INTERN;
- B. GRANT USAGE ON DATABASE SANDBOX TO ROLE INTERN;
- C. GRANT USAGE ON SCHEMA SANDBOX.PUBLIC TO ROLE INTERN;
- D. GRANT OWNERSHIP ON DATABASE SANDBOX TO USER INTERN;
- E. GRANT ALL PRIVILEGES ON DATABASE SANDBOX TO ROLE INTERN;

ANSWER: B C

QUESTION NO: 9

How do Snowflake databases that are created from shares differ from standard databases that are not created from shares? (Choose three.)

- A. Shared databases are read-only.
- B. Shared databases must be refreshed in order for new data to be visible.
- C. Shared databases cannot be cloned.
- D. Shared databases are not supported by Time Travel.
- E. Shared databases will have the PUBLIC or INFORMATION_SCHEMA schemas without explicitly granting these schemas to the share.
- F. Shared databases can also be created as transient databases.

ANSWER: A C E

QUESTION NO: 10

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.

ANSWER: C