

# Snowflake

## Exam Questions COF-C02

SnowPro Core Certification Exam (COF-C02)



### NEW QUESTION 1

- (Topic 1)

How long is Snowpipe data load history retained?

- A. As configured in the create pipe settings
- B. Until the pipe is dropped
- C. 64 days
- D. 14 days

**Answer:** C

#### **Explanation:**

Snowpipe data load history is retained for 64 days. This retention period allows users to review and audit the data load operations performed by Snowpipe over a significant period of time, which can be crucial for troubleshooting and ensuring data integrity.

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Snowflake Documentation on Snowpipe1

### NEW QUESTION 2

- (Topic 1)

Which of the following objects can be shared through secure data sharing?

- A. Masking policy
- B. Stored procedure
- C. Task
- D. External table

**Answer:** D

#### **Explanation:**

Secure data sharing in Snowflake allows users to share various objects between Snowflake accounts without physically copying the data, thus not consuming additional storage. Among the options provided, external tables can be shared through secure data sharing. External tables are used to query data directly from files in a stage without loading the data into Snowflake tables, making them suitable for sharing across different Snowflake accounts.

References:

- ? Snowflake Documentation on Secure Data Sharing
- ? SnowPro™ Core Certification Companion: Hands-on Preparation and Practice

### NEW QUESTION 3

- (Topic 1)

Which of the following Snowflake features provide continuous data protection automatically? (Select TWO).

- A. Internal stages
- B. Incremental backups
- C. Time Travel
- D. Zero-copy clones
- E. Fail-safe

**Answer:** CE

#### **Explanation:**

Snowflake's Continuous Data Protection (CDP) encompasses a set of features that help protect data stored in Snowflake against human error, malicious acts, and software failure. Time Travel allows users to access historical data (i.e., data that has been changed or deleted) for a defined period, enabling querying and restoring of data. Fail-safe is an additional layer of data protection that provides a recovery option in the event of significant data loss or corruption, which can only be performed by Snowflake. References:

- ? Continuous Data Protection | Snowflake Documentation1
- ? Data Storage Considerations | Snowflake Documentation2
- ? Snowflake SnowPro Core Certification Study Guide3
- ? Snowflake Data Cloud Glossary
- <https://docs.snowflake.com/en/user-guide/data-availability.html>

### NEW QUESTION 4

- (Topic 1)

What features does Snowflake Time Travel enable?

- A. Querying data-related objects that were created within the past 365 days
- B. Restoring data-related objects that have been deleted within the past 90 days
- C. Conducting point-in-time analysis for BI reporting
- D. Analyzing data usage/manipulation over all periods of time

**Answer:** BC

#### **Explanation:**

Snowflake Time Travel is a powerful feature that allows users to access historical data within a defined period. It enables two key capabilities:

? B. Restoring data-related objects that have been deleted within the past 90 days:

Time Travel can be used to restore tables, schemas, and databases that have been accidentally or intentionally deleted within the Time Travel retention period.

? C. Conducting point-in-time analysis for BI reporting: It allows users to query

historical data as it appeared at a specific point in time within the Time Travel retention period, which is crucial for business intelligence and reporting purposes.

While Time Travel does allow querying of past data, it is limited to the retention period set for the Snowflake account, which is typically 1 day for standard accounts

and can be extended up to 90 days for enterprise accounts. It does not enable querying or restoring objects created or deleted beyond the retention period, nor does it provide analysis over all periods of time.

References:

- ? Snowflake Documentation on Time Travel
- ? SnowPro® Core Certification Study Guide

#### NEW QUESTION 5

- (Topic 1)

Which Snowflake feature is used for both querying and restoring data?

- A. Cluster keys
- B. Time Travel
- C. Fail-safe
- D. Cloning

**Answer:** B

#### Explanation:

Snowflake's Time Travel feature is used for both querying historical data in tables and restoring and cloning historical data in databases, schemas, and tables. It allows users to access historical data within a defined period (1 day by default, up to 90 days for Snowflake Enterprise Edition) and is a key feature for data recovery and management. References: [COF-C02] SnowPro Core Certification Exam Study Guide

#### NEW QUESTION 6

- (Topic 1)

Which account usage views are used to evaluate the details of dynamic data masking? (Select TWO)

- A. ROLES
- B. POLICY\_REFERENCES
- C. QUERY\_HISTORY
- D. RESOURCE\_MONITOR
- E. ACCESS\_HISTORY

**Answer:** BE

#### Explanation:

To evaluate the details of dynamic data masking, the POLICY\_REFERENCES and ACCESS\_HISTORY views in the account\_usage schema are used. The POLICY\_REFERENCES view provides information about the objects to which a masking policy is applied, and the ACCESS\_HISTORY view contains details about access to the masked data, which can be used to audit and verify the application of dynamic data masking policies.

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Snowflake Documentation on Dynamic Data Masking1

#### NEW QUESTION 7

- (Topic 1)

Which of the following are valid methods for authenticating users for access into Snowflake? (Select THREE)

- A. SCIM
- B. Federated authentication
- C. TLS 1.2
- D. Key-pair authentication
- E. OAuth
- F. OCSP authentication

**Answer:** BDE

#### Explanation:

Snowflake supports several methods for authenticating users, including federated authentication, key-pair authentication, and OAuth. Federated authentication allows users to authenticate using their organization's identity provider. Key-pair authentication uses a public-private key pair for secure login, and OAuth is an open standard for access delegation commonly used for token-based authentication. References: Authentication policies | Snowflake Documentation, Authenticating to the server | Snowflake Documentation, External API authentication and secrets | Snowflake Documentation.

#### NEW QUESTION 8

- (Topic 1)

When is the result set cache no longer available? (Select TWO)

- A. When another warehouse is used to execute the query
- B. When another user executes the query
- C. When the underlying data has changed
- D. When the warehouse used to execute the query is suspended
- E. When it has been 24 hours since the last query

**Answer:** CE

#### Explanation:

The result set cache in Snowflake is invalidated and no longer available when the underlying data of the query results has changed, ensuring that queries return the most current data. Additionally, the cache expires after 24 hours to maintain the efficiency and accuracy of data retrieval.

### NEW QUESTION 9

- (Topic 1)

A user needs to create a materialized view in the schema MYDB.MYSCHEMA. Which statements will provide this access?

- A. GRANT ROLE MYROLE TO USER USER1;CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO ROLE MYROLE;
- B. GRANT ROLE MYROLE TO USER USER1;CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER USER1;
- C. GRANT ROLE MYROLE TO USER USER1;CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER1;
- D. GRANT ROLE MYROLE TO USER USER1;CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO MYROLE;

**Answer:** D

#### Explanation:

In Snowflake, to create a materialized view, the user must have the necessary privileges on the schema where the view will be created. These privileges are granted through roles, not directly to individual users. Therefore, the correct process is to grant the role to the user and then grant the privilege to create the materialized view to the role itself.

The statement GRANT ROLE MYROLE TO USER USER1; grants the specified role to the user, allowing them to assume that role and exercise its privileges. The subsequent statement CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO MYROLE; grants the privilege to create a materialized view within the specified schema to the role MYROLE. Any user who has been granted MYROLE can then create materialized views in MYDB.MYSCHEMA.

References:

? Snowflake Documentation on Roles

? Snowflake Documentation on Materialized Views

### NEW QUESTION 10

- (Topic 1)

A company strongly encourages all Snowflake users to self-enroll in Snowflake's default Multi-Factor Authentication (MFA) service to provide increased login security for users connecting to Snowflake.

Which application will the Snowflake users need to install on their devices in order to connect with MFA?

- A. Okta Verify
- B. Duo Mobile
- C. Microsoft Authenticator
- D. Google Authenticator

**Answer:** B

#### Explanation:

Snowflake's default Multi-Factor Authentication (MFA) service is powered by Duo Security. Users are required to install the Duo Mobile application on their devices to

use MFA for increased login security when connecting to Snowflake. This service is managed entirely by Snowflake, and users do not need to sign up separately with Duo1.

### NEW QUESTION 10

- (Topic 1)

A user has an application that writes a new Tile to a cloud storage location every 5 minutes.

What would be the MOST efficient way to get the files into Snowflake?

- A. Create a task that runs a copy into operation from an external stage every 5 minutes
- B. Create a task that puts the files in an internal stage and automate the data loading wizard
- C. Create a task that runs a GET operation to intermittently check for new files
- D. Set up cloud provider notifications on the Tile location and use Snowpipe with auto- ingest

**Answer:** D

#### Explanation:

The most efficient way to get files into Snowflake, especially when new files are being written to a cloud storage location at frequent intervals, is to use Snowpipe with auto-ingest. Snowpipe is Snowflake's continuous data ingestion service that loads data as soon as it becomes available in a cloud storage location. By setting up cloud provider notifications, Snowpipe can be triggered automatically whenever new files are written to the storage location, ensuring that the data is loaded into Snowflake with minimal latency and without the need for manual intervention or scheduling frequent tasks.

References:

? Snowflake Documentation on Snowpipe

? SnowPro® Core Certification Study Guide

### NEW QUESTION 11

- (Topic 1)

True or False: When you create a custom role, it is a best practice to immediately grant that role to ACCOUNTADMIN.

- A. True
- B. False

**Answer:** B

#### Explanation:

The ACCOUNTADMIN role is the most powerful role in Snowflake and should be limited to a select number of users within an organization. It is responsible for account-level configurations and should not be used for day-to-day object creation or management. Granting a custom role to ACCOUNTADMIN could inadvertently give broad access to users with this role, which is not a recommended security practice.

Reference:<https://docs.snowflake.com/en/user-guide/security-access-control-considerations.html>

### NEW QUESTION 12

- (Topic 1)

Which Snowflake objects track DML changes made to tables, like inserts, updates, and deletes?

- A. Pipes
- B. Streams
- C. Tasks
- D. Procedures

**Answer: B**

**Explanation:**

In Snowflake, Streams are the objects that track Data Manipulation Language (DML) changes made to tables, such as inserts, updates, and deletes. Streams record these changes along with metadata about each change, enabling actions to be taken using the changed data. This process is known as change data capture (CDC).

**NEW QUESTION 16**

- (Topic 1)

Which of the following Snowflake objects can be shared using a secure share? (Select TWO).

- A. Materialized views
- B. Sequences
- C. Procedures
- D. Tables
- E. Secure User Defined Functions (UDFs)

**Answer: DE**

**Explanation:**

Secure sharing in Snowflake allows users to share specific objects with other Snowflake accounts without physically copying the data, thus not consuming additional storage. Tables and Secure User Defined Functions (UDFs) are among the objects that can be shared using this feature. Materialized views, sequences, and procedures are not shareable objects in Snowflake.

References:

? [COF-C02] SnowPro Core Certification Exam Study Guide

? Snowflake Documentation on Secure Data Sharing1

**NEW QUESTION 17**

- (Topic 1)

How often are encryption keys automatically rotated by Snowflake?

- A. 30 Days
- B. 60 Days
- C. 90 Days
- D. 365 Days

**Answer: A**

**Explanation:**

Snowflake automatically rotates encryption keys when they are more than 30 days old. Active keys are retired, and new keys are created. This process is part of Snowflake's comprehensive security measures to ensure data protection and is managed entirely by the Snowflake service without requiring user intervention.

References:

? Understanding Encryption Key Management in Snowflake

**NEW QUESTION 19**

- (Topic 1)

What are two ways to create and manage Data Shares in Snowflake? (Choose two.)

- A. Via the Snowflake Web Interface (UI)
- B. Via thedata\_share=trueparameter
- C. Via SQL commands
- D. Via Virtual Warehouses

**Answer: AC**

**Explanation:**

In Snowflake, Data Shares can be created and managed in two primary ways:

? Via the Snowflake Web Interface (UI): Users can create and manage shares through the graphical interface provided by Snowflake, which allows for a user-friendly experience.

? Via SQL commands: Snowflake also allows the creation and management of

shares using SQL commands. This method is more suited for users who prefer scripting or need to automate the process.

Reference:<https://docs.snowflake.com/en/user-guide/data-sharing-provider.html>

**NEW QUESTION 22**

- (Topic 1)

A user unloaded a Snowflake table called mytable to an internal stage called mystage. Which command can be used to view the list of files that has been uploaded to the staged?

- A. list @mytable;
- B. list @%mytable;
- C. list @ %m.ystage;

D. list @mystage;

**Answer:** D

**Explanation:**

The command list @mystage; is used to view the list of files that have been uploaded to an internal stage in Snowflake. The list command displays the metadata for all files in the specified stage, which in this case is mystage. This command is particularly useful for verifying that files have been successfully unloaded from a Snowflake table to the stage and for managing the files within the stage.

References:

- ? Snowflake Documentation on Stages
- ? SnowPro® Core Certification Study Guide

**NEW QUESTION 27**

- (Topic 1)

The fail-safe retention period is how many days?

- A. 1 day
- B. 7 days
- C. 45 days
- D. 90 days

**Answer:** B

**Explanation:**

Fail-safe is a feature in Snowflake that provides an additional layer of data protection. After the Time Travel retention period ends, Fail-safe offers a non-configurable 7-day period during which historical data may be recoverable by Snowflake. This period is designed to protect against accidental data loss and is not intended for customer access. References: Understanding and viewing Fail-safe | Snowflake Documentation

**NEW QUESTION 29**

- (Topic 1)

How would you determine the size of the virtual warehouse used for a task?

- A. Root task may be executed concurrently (i.
- B. multiple instances), it is recommended to leave some margins in the execution window to avoid missing instances of execution
- C. Querying(select)the size of the stream content would help determine the warehouse siz
- D. For example, if querying large stream content, use a larger warehouse size
- E. If using the stored procedure to execute multiple SQL statements, it's best to test run the stored procedure separately to size the compute resource first
- F. Since task infrastructure is based on running the task body on schedule, it's recommended to configure the virtual warehouse for automatic concurrency handling using Multi-cluster warehouse (MCW) to match the task schedule

**Answer:** D

**Explanation:**

The size of the virtual warehouse for a task can be configured to handle concurrency automatically using a Multi-cluster warehouse (MCW). This is because tasks are designed to run their body on a schedule, and MCW allows for scaling compute resources to match the task's execution needs without manual intervention. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 33**

- (Topic 1)

In which use cases does Snowflake apply egress charges?

- A. Data sharing within a specific region
- B. Query result retrieval
- C. Database replication
- D. Loading data into Snowflake

**Answer:** C

**Explanation:**

Snowflake applies egress charges in the case of database replication when data is transferred out of a Snowflake region to another region or cloud provider. This is because the data transfer incurs costs associated with moving data across different networks. Egress charges are not applied for data sharing within the same region, query result retrieval, or loading data into Snowflake, as these actions do not involve data transfer across regions.

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Snowflake Documentation on Data Replication and Egress Charges1

**NEW QUESTION 38**

- (Topic 1)

Which command is used to unload data from a Snowflake table into a file in a stage?

- A. COPY INTO
- B. GET
- C. WRITE
- D. EXTRACT INTO

**Answer:** A

**Explanation:**

The COPY INTO command is used in Snowflake to unload data from a table into a file in a stage. This command allows for the export of data from Snowflake tables into flat files, which can then be used for further analysis, processing, or storage in external systems.

References:

? Snowflake Documentation on Unloading Data

? Snowflake SnowPro Core: Copy Into Command to Unload Rows to Files in Named Stage

#### NEW QUESTION 43

- (Topic 1)

Which of the following can be executed/called with Snowpipe?

- A. A User Defined Function (UDF)
- B. A stored procedure
- C. A single copy\_into statement
- D. A single insert into statement

**Answer:** C

#### Explanation:

Snowpipe is used for continuous, automated data loading into Snowflake. It uses a COPY INTO <table> statement within a pipe object to load data from files as soon as they are available in a stage. Snowpipe does not execute UDFs, stored procedures, or insert statements. References: Snowpipe | Snowflake Documentation

#### NEW QUESTION 47

- (Topic 1)

A user has unloaded data from Snowflake to a stage

Which SQL command should be used to validate which data was loaded into the stage?

- A. list @file stage
- B. show @file stage
- C. view @file stage
- D. verify @file stage

**Answer:** A

#### Explanation:

The list command in Snowflake is used to validate and display the list of files in a specified stage. When a user has unloaded data to a stage, running the list @file stage command will show all the files that have been uploaded to that stage, allowing the user to verify the data that was unloaded.

References:

? Snowflake Documentation on Stages

? SnowPro® Core Certification Study Guide

#### NEW QUESTION 49

- (Topic 1)

Which semi-structured file formats are supported when unloading data from a table? (Select TWO).

- A. ORC
- B. XML
- C. Avro
- D. Parquet
- E. JSON

**Answer:** DE

#### Explanation:

Semi-structured JSON, Parquet Snowflake supports unloading data in several semi-structured file formats, including Parquet and JSON. These formats allow for efficient storage and querying of semi-structured data, which can be loaded directly into Snowflake tables without requiring a predefined schema<sup>12</sup>.

[https://docs.snowflake.com/en/user-guide/data-unload-prepare.html#:~:text=Supported%20File%20Formats,-The%20following%20file&text=Delimited%20\(CSV%2C%20TSV%2C%20etc.\)](https://docs.snowflake.com/en/user-guide/data-unload-prepare.html#:~:text=Supported%20File%20Formats,-The%20following%20file&text=Delimited%20(CSV%2C%20TSV%2C%20etc.))

#### NEW QUESTION 54

- (Topic 1)

What data is stored in the Snowflake storage layer? (Select TWO).

- A. Snowflake parameters
- B. Micro-partitions
- C. Query history
- D. Persisted query results
- E. Standard and secure view results

**Answer:** BD

#### Explanation:

The Snowflake storage layer is responsible for storing data in an optimized, compressed, columnar format. This includes micro-partitions, which are the fundamental storage units that contain the actual data stored in Snowflake. Additionally, persisted query results, which are the results of queries that have been materialized and stored for future use, are also kept within this layer. This design allows for efficient data retrieval and management within the Snowflake architecture<sup>1</sup>.

References:

? [COF-C02] SnowPro Core Certification Exam Study Guide

? Key Concepts & Architecture | Snowflake Documentation2

### NEW QUESTION 58

- (Topic 1)

What Snowflake role must be granted for a user to create and manage accounts?

- A. ACCOUNTADMIN
- B. ORGADMIN
- C. SECURITYADMIN
- D. SYSADMIN

**Answer:** A

#### Explanation:

The ACCOUNTADMIN role is required for a user to create and manage accounts in Snowflake. This role has the highest level of privileges and is responsible for managing all aspects of the Snowflake account, including the ability to create and manage other user accounts<sup>1</sup>.

<https://docs.snowflake.com/en/user-guide/security-access-control-considerations.html>

### NEW QUESTION 63

- (Topic 1)

What happens when an external or an internal stage is dropped? (Select TWO).

- A. When dropping an external stage, the files are not removed and only the stage is dropped
- B. When dropping an external stage, both the stage and the files within the stage are removed
- C. When dropping an internal stage, the files are deleted with the stage and the files are recoverable
- D. When dropping an internal stage, the files are deleted with the stage and the files are not recoverable
- E. When dropping an internal stage, only selected files are deleted with the stage and are not recoverable

**Answer:** AD

#### Explanation:

When an external stage is dropped in Snowflake, the reference to the external storage location is removed, but the actual files within the external storage (like Amazon S3, Google Cloud Storage, or Microsoft Azure) are not deleted. This means that the data remains intact in the external storage location, and only the stage object in Snowflake is removed.

On the other hand, when an internal stage is dropped, any files that were uploaded to the stage are deleted along with the stage itself. These files are not recoverable once the internal stage is dropped, as they are permanently removed from Snowflake's storage. References:

? [COF-C02] SnowPro Core Certification Exam Study Guide

? Snowflake Documentation on Stages

### NEW QUESTION 67

- (Topic 2)

What occurs when a pipe is recreated using the CREATE OR REPLACE PIPE command?

- A. The Pipe load history is reset to empty.
- B. The REFRESH command is executed.
- C. The stage will be purged.
- D. The destination table is truncated.

**Answer:** A

#### Explanation:

When a pipe is recreated using the CREATE OR REPLACE

PIPE command, the load history of the pipe is reset. This means that Snowpipe will consider all files in the stage as new and will attempt to load them, even if they were loaded previously by the old pipe<sup>2</sup>.

### NEW QUESTION 68

- (Topic 2)

How does Snowflake Fail-safe protect data in a permanent table?

- A. Fail-safe makes data available up to 1 day, recoverable by user operations.
- B. Fail-safe makes data available for 7 days, recoverable by user operations.
- C. Fail-safe makes data available for 7 days, recoverable only by Snowflake Support.
- D. Fail-safe makes data available up to 1 day, recoverable only by Snowflake Support.

**Answer:** C

#### Explanation:

Snowflake's Fail-safe provides a 7-day period during which data in a permanent table may be recoverable, but only by Snowflake Support, not by user operations<sup>3</sup>.

### NEW QUESTION 69

- (Topic 2)

What are the correct parameters for time travel and fail-safe in the Snowflake Enterprise Edition?

- A. Default Time Travel Retention is set to 0 day
- B. Maximum Time Travel Retention is 30 day
- C. Fail Safe retention time is 1 day.
- D. Default Time Travel Retention is set to 1 da

- E. Maximum Time Travel Retention is 365 day
- F. Fail Safe retention time is 7 days.
- G. Default Time Travel Retention is set to 0 day
- H. Maximum Time Travel Retention is 90 day
- I. Fail Safe retention time is 7 days.
- J. Default Time Travel Retention is set to 1 da
- K. Maximum Time Travel Retention is 90 day
- L. Fail Safe retention time is 7 days.
- M. Default Time Travel Retention is set to 7 day
- N. Maximum Time Travel Retention is 1 da
- O. Fail Safe retention time is 90 days.
- P. Default Time Travel Retention is set to 90 day
- Q. Maximum Time Travel Retention is 7 day
- R. Fail Safe retention time is 356 days.

**Answer:** D

**Explanation:**

In the Snowflake Enterprise Edition, the default Time Travel retention is set to 1 day, the maximum Time Travel retention can be set up to 90 days, and the Fail-safe retention time is 7 days<sup>3</sup>.

**NEW QUESTION 74**

- (Topic 2)

The is the minimum Fail-safe retention time period for transient tables?

- A. 1 day
- B. 7 days
- C. 12 hours
- D. 0 days

**Answer:** D

**Explanation:**

Transient tables in Snowflake have a minimum Fail-safe retention time period of 0 days. This means that once the Time Travel retention period ends, there is no additional Fail-safe period for transient tables

**NEW QUESTION 78**

- (Topic 2)

Which of the following statements apply to Snowflake in terms of security? (Choose two.)

- A. Snowflake leverages a Role-Based Access Control (RBAC) model.
- B. Snowflake requires a user to configure an IAM user to connect to the database.
- C. All data in Snowflake is encrypted.
- D. Snowflake can run within a user's own Virtual Private Cloud (VPC).
- E. All data in Snowflake is compressed.

**Answer:** AC

**Explanation:**

Snowflake uses a Role-Based Access Control (RBAC) model to manage access to data and resources. Additionally, Snowflake ensures that all data is encrypted, both at rest and in transit, to provide a high level of security for data stored within the platform. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 79**

- (Topic 2)

The Snowflake cloud services layer is responsible for which tasks? (Choose two.)

- A. Local disk caching
- B. Authentication and access control
- C. Metadata management
- D. Query processing
- E. Database storage

**Answer:** BC

**Explanation:**

The Snowflake cloud services layer is responsible for tasks such as authentication and access control, ensuring secure access to the platform, and metadata management, which involves organizing and maintaining information about the data stored in Snowflake<sup>56</sup>.

**NEW QUESTION 84**

- (Topic 2)

In a Snowflake role hierarchy, what is the top-level role?

- A. SYSADMIN
- B. ORGADMIN
- C. ACCOUNTADMIN
- D. SECURITYADMIN

**Answer:** C

**Explanation:**

In a Snowflake role hierarchy, the top-level role is ACCOUNTADMIN. This role has the highest level of privileges and is capable of performing all administrative functions within the Snowflake account

**NEW QUESTION 85**

- (Topic 2)

Which Snowflake layer is always leveraged when accessing a query from the result cache?

- A. Metadata
- B. Data Storage
- C. Compute
- D. Cloud Services

**Answer:** D

**Explanation:**

The Cloud Services layer in Snowflake is responsible for managing the result cache. When a query is executed, the results are stored in this cache, and subsequent identical queries can leverage these cached results without re-executing the entire query.

**NEW QUESTION 90**

- (Topic 2)

Which of the following statements describe features of Snowflake data caching? (Choose two.)

- A. When a virtual warehouse is suspended, the data cache is saved on the remote storage layer.
- B. When the data cache is full, the least-recently used data will be cleared to make room.
- C. A user can only access their own queries from the query result cache.
- D. A user must set USE\_METADATA\_CACHE to TRUE to use the metadata cache in queries.
- E. The RESULT\_SCAN table function can access and filter the contents of the query result cache.

**Answer:** BE

**Explanation:**

Snowflake's data caching features include the ability to clear the least-recently used data when the data cache is full to make room for new data. Additionally, the RESULT\_SCAN table function can access and filter the contents of the query result cache, allowing users to retrieve and work with the results of previous queries. The other statements are incorrect: the data cache is not saved on the remote storage layer when a virtual warehouse is suspended, users can access queries from the query result cache that were run by other users, and there is no setting called USE\_METADATA\_CACHE in Snowflake. References: Caching in the Snowflake Cloud Data Platform, Optimizing the warehouse cache

**NEW QUESTION 93**

- (Topic 2)

Which of the following are best practices for loading data into Snowflake? (Choose three.)

- A. Aim to produce data files that are between 100 MB and 250 MB in size, compressed.
- B. Load data from files in a cloud storage service in a different region or cloud platform from the service or region containing the Snowflake account, to save on cost.
- C. Enclose fields that contain delimiter characters in single or double quotes.
- D. Split large files into a greater number of smaller files to distribute the load among the compute resources in an active warehouse.
- E. When planning which warehouse(s) to use for data loading, start with the largest warehouse possible.
- F. Partition the staged data into large folders with random paths, allowing Snowflake to determine the best way to load each file.

**Answer:** ACD

**Explanation:**

Best practices for loading data into Snowflake include aiming for data file sizes between 100 MB and 250 MB when compressed, as this size is optimal for parallel processing and minimizes overhead. Enclosing fields with delimiter characters in quotes ensures proper field recognition during the load process. Splitting large files into smaller ones allows for better distribution of the load across compute resources, enhancing performance and efficiency.

**NEW QUESTION 98**

- (Topic 2)

A company needs to allow some users to see Personally Identifiable Information (PII) while limiting other users from seeing the full value of the PII. Which Snowflake feature will support this?

- A. Row access policies
- B. Data masking policies
- C. Data encryption
- D. Role based access control

**Answer:** B

**Explanation:**

Data masking policies in Snowflake allow for the obfuscation of specific data within a field, enabling some users to see the full data while limiting others. This feature is particularly useful for handling PII, ensuring that sensitive information is only visible to authorized users.

**NEW QUESTION 100**

- (Topic 2)

The Snowflake Search Optimization Services supports improved performance of which kind of query?

- A. Queries against large tables where frequent DML occurs
- B. Queries against tables larger than 1 TB
- C. Selective point lookup queries
- D. Queries against a subset of columns in a table

**Answer:** C

**Explanation:**

The Snowflake Search Optimization Service is designed to support improved performance for selective point lookup queries. These are queries that retrieve specific records from a database, often based on a unique identifier or a small set of criteria<sup>3</sup>.

**NEW QUESTION 101**

- (Topic 2)

Which methods can be used to delete staged files from a Snowflake stage? (Choose two.)

- A. Use the DROP <file> command after the load completes.
- B. Specify the TEMPORARY option when creating the file format.
- C. Specify the PURGE copy option in the COPY INTO <table> command.
- D. Use the REMOVE command after the load completes.
- E. Use the DELETE LOAD HISTORY command after the load completes.

**Answer:** CD

**Explanation:**

To delete staged files from a Snowflake stage, you can specify the PURGE option in the COPY INTO <table> command, which will automatically delete the files after they have been successfully loaded. Additionally, you can use the REMOVE command after the load completes to manually delete the files from the stage<sup>12</sup>.  
References = DROP STAGE, REMOVE

**NEW QUESTION 103**

- (Topic 2)

Which of the following is a data tokenization integration partner?

- A. Protegrity
- B. Tableau
- C. DBeaver
- D. SAP

**Answer:** A

**Explanation:**

Protegrity is listed as a data tokenization integration partner for Snowflake. This partnership allows Snowflake users to utilize Protegrity's tokenization solutions within the Snowflake environment<sup>3</sup>.  
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

**NEW QUESTION 106**

- (Topic 2)

A single user of a virtual warehouse has set the warehouse to auto-resume and auto-suspend after 10 minutes. The warehouse is currently suspended and the user performs the following actions:

- \* 1. Runs a query that takes 3 minutes to complete
  - \* 2. Leaves for 15 minutes
  - \* 3. Returns and runs a query that takes 10 seconds to complete
  - \* 4. Manually suspends the warehouse as soon as the last query was completed
- When the user returns, how much billable compute time will have been consumed?

- A. 4 minutes
- B. 10 minutes
- C. 14 minutes
- D. 24 minutes

**Answer:** C

**Explanation:**

The billable compute time includes the time the warehouse is running queries plus the auto-suspend time after the last query if the warehouse is not manually suspended. In this scenario, the warehouse runs for 3 minutes, suspends after 10 minutes of inactivity, resumes for a 10-second query, and then is manually suspended. The total billable time is the sum of the initial 3 minutes, the 10 minutes of auto-suspend time, and the brief period for the 10-second query, which is rounded up to the next full minute due to Snowflake's billing practices. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 107**

- (Topic 2)

By default, which Snowflake role is required to create a share?

- A. ORGADMIN
- B. SECURITYADMIN
- C. SHAREADMIN
- D. ACCOUNTADMIN

**Answer:** D

**Explanation:**

By default, the Snowflake role required to create a share is ACCOUNTADMIN (D). This role has the necessary privileges to perform administrative tasks, including creating shares for data sharing purposes

**NEW QUESTION 109**

- (Topic 2)

What are best practice recommendations for using the ACCOUNTADMIN system-defined role in Snowflake? (Choose two.)

- A. Ensure all ACCOUNTADMIN roles use Multi-factor Authentication (MFA).
- B. All users granted ACCOUNTADMIN role must be owned by the ACCOUNTADMIN role.
- C. The ACCOUNTADMIN role must be granted to only one user.
- D. Assign the ACCOUNTADMIN role to at least two users, but as few as possible.
- E. All users granted ACCOUNTADMIN role must also be granted SECURITYADMIN role.

**Answer:** AD

**Explanation:**

Best practices for using the ACCOUNTADMIN role include ensuring that all users with this role use Multi-factor Authentication (MFA) for added security. Additionally, it is recommended to assign the ACCOUNTADMIN role to at least two users to avoid delays in case of password recovery issues, but to as few users as possible to maintain strict control over account-level operations.

**NEW QUESTION 110**

- (Topic 2)

If 3 size Small virtual warehouse is made up of two servers, how many servers make up a Large warehouse?

- A. 4
- B. 8
- C. 16
- D. 32

**Answer:** B

**Explanation:**

In Snowflake, each size increase in virtual warehouses doubles the number of servers. Therefore, if a size Small virtual warehouse is made up of two servers, a Large warehouse, which is two sizes larger, would be made up of eight servers (2 servers for Small, 4 for Medium, and 8 for Large).

Size specifies the amount of compute resources available per cluster in a warehouse. Snowflake supports the following warehouse sizes:

Warehouse Size	Credits / Hour	Credits / Second	Notes
X-Small	1	0.0003	Default size for warehouses created using CREATE WAREHOUSE.
Small	2	0.0006	
Medium	4	0.0011	
Large	8	0.0022	
X-Large	16	0.0044	Default for warehouses created in the web interface.
2X-Large	32	0.0089	
3X-Large	64	0.0178	
4X-Large	128	0.0356	
5X-Large	256	0.0711	Preview feature.
6X-Large	512	0.1422	Preview feature.

<https://docs.snowflake.com/en/user-guide/warehouses-overview.html>

**NEW QUESTION 114**

- (Topic 2)

What happens to historical data when the retention period for an object ends?

- A. The data is cloned into a historical object.
- B. The data moves to Fail-safe
- C. Time Travel on the historical data is dropped.
- D. The object containing the historical data is dropped.

**Answer:** C

**Explanation:**

When the retention period for an object ends in Snowflake, Time Travel on the historical data is dropped ©. This means that the ability to access historical data via Time Travel is no longer available once the retention period has expired.

#### NEW QUESTION 115

- (Topic 2)

Which of the following describes the Snowflake Cloud Services layer?

- A. Coordinates activities in the Snowflake account
- B. Executes queries submitted by the Snowflake account users
- C. Manages quotas on the Snowflake account storage
- D. Manages the virtual warehouse cache to speed up queries

**Answer:** A

#### Explanation:

The Snowflake Cloud Services layer is a collection of services that coordinate activities across Snowflake, tying together all the different components to process user requests, from login to query dispatch<sup>1</sup>.

References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation<sup>1</sup>

#### NEW QUESTION 120

- (Topic 2)

A user created a new worksheet within the Snowsight UI and wants to share this with teammates

How can this worksheet be shared?

- A. Create a zero-copy clone of the worksheet and grant permissions to teammates
- B. Create a private Data Exchange so that any teammate can use the worksheet
- C. Share the worksheet with teammates within Snowsight
- D. Create a database and grant all permissions to teammates

**Answer:** C

#### Explanation:

Worksheets in Snowsight can be shared directly with other Snowflake users within the same account. This feature allows for collaboration and sharing of SQL queries or Python code, as well as other data manipulation tasks<sup>1</sup>.

#### NEW QUESTION 124

- (Topic 2)

True or False: Snowpipe via REST API can only reference External Stages as source.

- A. True
- B. False

**Answer:** B

#### Explanation:

Snowpipe via REST API can reference both named internal stages within Snowflake and external stages, such as Amazon S3, Google Cloud Storage, or Microsoft Azure<sup>1</sup>. This means that Snowpipe is not limited to only external stages as a source for data loading.

References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation<sup>1</sup>

Reference: <https://community.snowflake.com/s/article/Making-Transient-table-by-Default>

#### NEW QUESTION 126

- (Topic 2)

Which statements are true concerning Snowflake's underlying cloud infrastructure? (Select THREE),

- A. Snowflake data and services are deployed in a single availability zone within a cloud provider's region.
- B. Snowflake data and services are available in a single cloud provider and a single region, the use of multiple cloud providers is not supported.
- C. Snowflake can be deployed in a customer's private cloud using the customer's own compute and storage resources for Snowflake compute and storage
- D. Snowflake uses the core compute and storage services of each cloud provider for its own compute and storage
- E. All three layers of Snowflake's architecture (storage, compute, and cloud services) are deployed and managed entirely on a selected cloud platform
- F. Snowflake data and services are deployed in at least three availability zones within a cloud provider's region

**Answer:** DEF

#### Explanation:

Snowflake's architecture is designed to operate entirely on cloud infrastructure. It uses the core compute and storage services of each cloud provider, which allows it to leverage the scalability and reliability of cloud resources. Snowflake's services are deployed across multiple availability zones within a cloud provider's region to ensure high availability and fault tolerance. References: [COF-C02] SnowPro Core Certification Exam Study Guide

#### NEW QUESTION 128

- (Topic 2)

What are the responsibilities of Snowflake's Cloud Service layer? (Choose three.)

- A. Authentication
- B. Resource management
- C. Virtual warehouse caching
- D. Query parsing and optimization
- E. Query execution
- F. Physical storage of micro-partitions

**Answer:** ABD

**Explanation:**

The responsibilities of Snowflake's Cloud Service layer include authentication (A), which ensures secure access to the platform; resource management (B), which involves allocating and managing compute resources; and query parsing and optimization (D), which improves the efficiency and performance of SQL query execution.

**NEW QUESTION 133**

- (Topic 2)

Why does Snowflake recommend file sizes of 100-250 MB compressed when loading data?

- A. Optimizes the virtual warehouse size and multi-cluster setting to economy mode
- B. Allows a user to import the files in a sequential order
- C. Increases the latency staging and accuracy when loading the data
- D. Allows optimization of parallel operations

**Answer:** D

**Explanation:**

Snowflake recommends file sizes between 100-250 MB compressed when loading data to optimize parallel processing. Smaller, compressed files can be loaded in parallel, which maximizes the efficiency of the virtual warehouses and speeds up the data loading process.

**NEW QUESTION 134**

- (Topic 2)

A user has unloaded data from a Snowflake table to an external stage.

Which command can be used to verify if data has been uploaded to the external stage named my\_stage?

- A. view @my\_stage
- B. list @my\_stage
- C. show @my\_stage
- D. display @my\_stage

**Answer:** B

**Explanation:**

The list @my\_stage command in Snowflake can be used to verify if data has been uploaded to an external stage named my\_stage. This command provides a list of files that are present in the specified stage.

**NEW QUESTION 138**

- (Topic 2)

Which services does the Snowflake Cloud Services layer manage? (Choose two.)

- A. Compute resources
- B. Query execution
- C. Authentication
- D. Data storage
- E. Metadata

**Answer:** CE

**Explanation:**

The Snowflake Cloud Services layer manages various services, including authentication and metadata management. This layer ties together all the different components of Snowflake to process user requests, manage sessions, and control access.

**NEW QUESTION 140**

- (Topic 2)

What do the terms scale up and scale out refer to in Snowflake? (Choose two.)

- A. Scaling out adds clusters of the same size to a virtual warehouse to handle more concurrent queries.
- B. Scaling out adds clusters of varying sizes to a virtual warehouse.
- C. Scaling out adds additional database servers to an existing running cluster to handle more concurrent queries.
- D. Snowflake recommends using both scaling up and scaling out to handle more concurrent queries.
- E. Scaling up resizes a virtual warehouse so it can handle more complex workloads.
- F. Scaling up adds additional database servers to an existing running cluster to handle larger workloads.

**Answer:** AE

**Explanation:**

Scaling out in Snowflake involves adding clusters of the same size to a virtual warehouse, which allows for handling more concurrent queries without affecting the performance of individual queries. Scaling up refers to resizing a virtual warehouse to increase its compute resources, enabling it to handle more complex workloads and larger queries more efficiently.

**NEW QUESTION 144**

- (Topic 3)

How does Snowflake handle the bulk unloading of data into single or multiple files?

- A. It assigns each unloaded data file a unique name.
- B. It uses the put command to download the data by default.
- C. It uses COPY INTO <location> for bulk unloading where the default option is SINGLE - TRUE.

D. It uses COPY INTO <location> to copy the data from a table into one or more files in an external stage only.

**Answer:** A

**Explanation:**

When unloading data, Snowflake assigns each file a unique name to ensure there is no overlap or confusion between files. This is part of the bulk unloading process where data is exported from Snowflake tables into flat files.

**NEW QUESTION 147**

- (Topic 3)

Using variables in Snowflake is denoted by using which SQL character?

- A. @
- B. &
- C. \$
- D. #

**Answer:** C

**Explanation:**

VeryComprehensiveExplanation=InSnowflake,variablesaredenotedbyadollarsign().Variable s can be used in SQL statements where a literal constant is allowed, and they must be prefixed with a \$ sign to distinguish them from bind values and column names.

**NEW QUESTION 152**

- (Topic 3)

What internal stages are available in Snowflake? (Choose three.)

- A. Schema stage
- B. Named stage
- C. User stage
- D. Stream stage
- E. Table stage
- F. Database stage

**Answer:** BCE

**Explanation:**

Snowflake supports three types of internal stages: Named, User, and Table stages. These stages are used for staging data files to be loaded into Snowflake tables. Schema, Stream, and Database stages are not supported as internal stages in Snowflake. References: Snowflake Documentation1.

**NEW QUESTION 154**

- (Topic 3)

What column type does a Kafka connector store formatted information in a single column?

- A. ARRAY
- B. OBJECT
- C. VARCHAR
- D. VARIANT

**Answer:** D

**Explanation:**

The Kafka connector stores formatted information in a single column of type VARIANT. This column type is used to store semi-structured data like JSON or Avro, which allows for flexibility in the data structure

**NEW QUESTION 157**

- (Topic 3)

For the ALLOWED VALUES tag property, what is the MAXIMUM number of possible string values for a single tag?

- A. 10
- B. 50
- C. 64
- D. 256

**Answer:** D

**Explanation:**

For the ALLOWED VALUES tag property, the maximum number of possible string values for a single tag is 256. This allows for a wide range of values to be assigned to a tag when it is set on an object

**NEW QUESTION 159**

- (Topic 3)

Which of the following practices are recommended when creating a user in Snowflake? (Choose two.)

- A. Configure the user to be initially disabled.
- B. Force an immediate password change.
- C. Set a default role for the user.

- D. Set the number of minutes to unlock to 15 minutes.
- E. Set the user's access to expire within a specified timeframe.

**Answer:** BC

#### NEW QUESTION 160

- (Topic 3)

If a virtual warehouse runs for 61 seconds, shuts down, and then restarts and runs for 30 seconds, for how many seconds is it billed?

- A. 60
- B. 91
- C. 120
- D. 121

**Answer:** D

#### Explanation:

Snowflake's billing for virtual warehouses is per-second, with a minimum of 60 seconds for each time the warehouse is started or resumed. Therefore, if a warehouse runs for 61 seconds, it is billed for 61 seconds. If it is then shut down and restarted, running for an additional 30 seconds, it is billed for another 60 seconds (the minimum charge for a restart), totaling 121 seconds.

#### NEW QUESTION 165

- (Topic 3)

Which data type can store more than one type of data structure?

- A. JSON
- B. BINARY
- C. VARCHAR
- D. VARIANT

**Answer:** D

#### Explanation:

The VARIANT data type in Snowflake can store multiple types of data structures, as it is designed to hold semi-structured data. It can contain any other data type, including OBJECT and ARRAY, which allows it to represent various data structures.

#### NEW QUESTION 168

- (Topic 3)

How long does Snowflake retain information in the ACCESS HISTORY view?

- A. 7 days
- B. 14 days
- C. 28 days
- D. 365 days

**Answer:** D

#### Explanation:

Snowflake retains information in the ACCESS HISTORY view for 365 days. This allows users to query the access history of Snowflake objects within the last year.

#### NEW QUESTION 169

- (Topic 3)

Which query contains a Snowflake hosted file URL in a directory table for a stage named bronzestage?

- A. list @bronzestage;
- B. select \* from directory(@bronzestage);
- C. select metadata\$filename from @bronzestage;
- D. select \* from table(information\_schema.stage\_directory\_file\_registration\_history( stage name=>'bronzestage1'));

**Answer:** B

#### Explanation:

The query that contains a Snowflake hosted file URL in a directory table for a stage named bronzestage is select \* from directory(@bronzestage). This query retrieves a list of all files on the stage along with metadata, including the Snowflake file URL for each file.

#### NEW QUESTION 171

- (Topic 3)

Which clients does Snowflake support Multi-Factor Authentication (MFA) token caching for? (Select TWO).

- A. GO driver
- B. Node.js driver
- C. ODBC driver
- D. Python connector
- E. Spark connector

**Answer:** CD

**Explanation:**

Multi-Factor Authentication (MFA) token caching is typically supported for clients that maintain a persistent connection or session with Snowflake, such as the ODBC driver and Python connector, to reduce the need for repeated MFA challenges. References: Based on general security practices in cloud services as of 2021.

**NEW QUESTION 176**

- (Topic 3)

Which transformation is supported by a COPY INTO <table> command?

- A. Filter using a where clause
- B. Filter using a limit keyword
- C. Cast using a SELECT statement
- D. Order using an ORDER BY clause

**Answer: C**

**Explanation:**

The COPY INTO <table> command in Snowflake supports transformations such as casting using a SELECT statement. This allows for the transformation of data types as the data is being loaded into the table, which can be particularly useful when the data types in the source files do not match the data types in the target table

**NEW QUESTION 178**

- (Topic 3)

How would a user execute a series of SQL statements using a task?

- A. Include the SQL statements in the body of the task CREATE TASK mytask .. AS INSERT INTO target1 SELECT .. FROM stream\_s1 WHERE .. INSERT INTO target2 SELECT .. FROM stream\_s1WHERE ..
- B. A stored procedure can have only one DML statement per stored procedure invocation and therefore the user should sequence stored procedure calls in the taskdefinition CREATE TASK mytask .... AS call stored\_proc1(); call stored\_proc2();
- C. Use a stored procedure executing multiple SQL statements and invoke the stored procedure from the task
- D. CREATE TASK mytask .... AS call stored\_proc\_multiple\_statements\_inside();
- E. Create a task for each SQL statement (e.
- F. resulting in task1, task2, etc.) and string the series of SQL statements by having a control task calling task1, task2, et
- G. sequentially.

**Answer: C**

**Explanation:**

To execute a series of SQL statements using a task, a user would use a stored procedure that contains multiple SQL statements and invoke this stored procedure from the task. References: Snowflake Documentation2.

**NEW QUESTION 182**

- (Topic 3)

What privilege should a user be granted to change permissions for new objects in a managed access schema?

- A. Grant the OWNERSHIP privilege on the schema.
- B. Grant the OWNERSHIP privilege on the database.
- C. Grant the MANAGE GRANTS global privilege.
- D. Grant ALL privileges on the schema.

**Answer: C**

**Explanation:**

To change permissions for new objects in a managed access schema, a user should be granted the MANAGE GRANTS global privilege. This privilege allows the user to manage access control through grants on all securable objects within Snowflake2. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 183**

- (Topic 3)

Which stream type can be used for tracking the records in external tables?

- A. Append-only
- B. External
- C. Insert-only
- D. Standard

**Answer: B**

**Explanation:**

The stream type that can be used for tracking the records in external tables is ??External??. This type of stream is specifically designed to track changes in external tables

**NEW QUESTION 188**

- (Topic 3)

Which of the following statements describes a schema in Snowflake?

- A. A logical grouping of objects that belongs to a single database
- B. A logical grouping of objects that belongs to multiple databases

- C. A named Snowflake object that includes all the information required to share a database
- D. A uniquely identified Snowflake account within a business entity

**Answer:** A

**Explanation:**

A schema in Snowflake is a logical grouping of database objects, such as tables and views, that belongs to a single database. Each schema is part of a namespace in Snowflake, which is inferred from the current database and schema in use for the session<sup>5</sup>

**NEW QUESTION 191**

- (Topic 3)

The bulk data load history that is available upon completion of the COPY statement is stored where and for how long?

- A. In the metadata of the target table for 14 days
- B. In the metadata of the pipe for 14 days
- C. In the metadata of the target table for 64 days
- D. In the metadata of the pipe for 64 days

**Answer:** D

**Explanation:**

The bulk data load history available after a COPY statement is stored in the metadata of the pipe and is retained for 64 days<sup>1</sup>.

**NEW QUESTION 192**

- (Topic 3)

What can a Snowflake user do in the Activity section in Snowsight?

- A. Create dashboards.
- B. Write and run SQL queries.
- C. Explore databases and objects.
- D. Explore executed query performance.

**Answer:** D

**Explanation:**

In the Activity section in Snowsight, Snowflake users can explore the performance of executed queries. This includes monitoring queries, viewing details about queries, including performance data, and exploring each step of an executed query in the query profile<sup>1</sup>.

**NEW QUESTION 194**

- (Topic 3)

What action can a user take to address query concurrency issues?

- A. Enable the query acceleration service.
- B. Enable the search optimization service.
- C. Add additional clusters to the virtual warehouse
- D. Resize the virtual warehouse to a larger instance size.

**Answer:** C

**Explanation:**

To address query concurrency issues, a user can add additional clusters to the virtual warehouse. This allows for the distribution of queries across multiple clusters, reducing the load on any single cluster and improving overall query performance<sup>2</sup>.

**NEW QUESTION 198**

- (Topic 3)

Which kind of Snowflake table stores file-level metadata for each file in a stage?

- A. Directory
- B. External
- C. Temporary
- D. Transient

**Answer:** A

**Explanation:**

The kind of Snowflake table that stores file-level metadata for each file in a stage is a directory table. A directory table is an implicit object layered on a stage and stores file-level metadata about the data files in the stage<sup>3</sup>.

**NEW QUESTION 199**

- (Topic 3)

What are benefits of using Snowpark with Snowflake? (Select TWO).

- A. Snowpark uses a Spark engine to generate optimized SQL query plans.
- B. Snowpark automatically sets up Spark within Snowflake virtual warehouses.
- C. Snowpark does not require that a separate cluster be running outside of Snowflake.
- D. Snowpark allows users to run existing Spark code on virtual warehouses without the need to reconfigure the code.
- E. Snowpark executes as much work as possible in the source databases for all operations including User-Defined Functions (UDFs).

**Answer:** CD

**Explanation:**

Snowpark is designed to bring the data programmability to Snowflake, enabling developers to write code in familiar languages like Scala, Java, and Python. It allows for the execution of these codes directly within Snowflake's virtual warehouses, eliminating the need for a separate cluster. Additionally, Snowpark's compatibility with Spark allows users to leverage their existing Spark code with minimal changes<sup>1</sup>.

**NEW QUESTION 202**

- (Topic 3)

What can a Snowflake user do in the Admin area of Snowsight?

- A. Analyze query performance.
- B. Write queries and execute them.
- C. Provide an overview of the listings in the Snowflake Marketplace.
- D. Connect to Snowflake partners to explore extended functionality.

**Answer:** A

**Explanation:**

In the Admin area of Snowsight, users can analyze query performance, manage Snowflake warehouses, set up and view details about resource monitors, manage users and roles, and administer Snowflake accounts in their organization<sup>2</sup>.

**NEW QUESTION 204**

- (Topic 3)

A data provider wants to share data with a consumer who does not have a Snowflake account. The provider creates a reader account for the consumer following these steps:

- \* 1. Created a user called "CONSUMER"
- \* 2. Created a database to hold the share and an extra-small warehouse to query the data
- \* 3. Granted the role PUBLIC the following privileges: Usage on the warehouse, database, and schema, and SELECT on all the objects in the share

Based on this configuration what is true of the reader account?

- A. The reader account will automatically use the Standard edition of Snowflake.
- B. The reader account compute will be billed to the provider account.
- C. The reader account can clone data the provider has shared, but cannot re-share it.
- D. The reader account can create a copy of the shared data using CREATE TABLE AS...

**Answer:** B

**Explanation:**

The reader account compute will be billed to the provider account.

Very Comprehensive Explanation

In Snowflake, when a provider creates a reader account for a consumer who does not have a Snowflake account, the compute resources used by the reader account are billed to the provider's account. This allows the consumer to query the shared data without incurring any costs. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 208**

- (Topic 3)

Which of the following describes the Snowflake Cloud Services layer?

- A. Coordinates activities in the Snowflake account
- B. Executes queries submitted by the Snowflake account users
- C. Manages quotas on the Snowflake account storage
- D. Manages the virtual warehouse cache to speed up queries

**Answer:** A

**Explanation:**

The Snowflake Cloud Services layer coordinates activities within the Snowflake account. It is responsible for tasks such as authentication, infrastructure management, metadata management, query parsing and optimization, and access control. References: Based on general cloud database architecture knowledge.

**NEW QUESTION 209**

- (Topic 3)

Which Snowflake object helps evaluate virtual warehouse performance impacted by query queuing?

- A. Resource monitor
- B. Account\_usag
- C. query\_history
- D. Information\_schema.warehouse\_load\_history
- E. Information schema.warehouse metering history

**Answer:** C

**Explanation:**

The Snowflake object that helps evaluate virtual warehouse performance impacted by query queuing is the Information\_schema.warehouse\_load\_history. This view provides historical data about the load on a warehouse, including the average number of queries that were running or queued within a specific interval, which can be used to assess performance and identify potential issues with query queuing<sup>3</sup>.

#### NEW QUESTION 213

- (Topic 3)

Which Snowflake URL type is used by directory tables?

- A. File
- B. Pre-signed
- C. Scoped
- D. Virtual-hosted style

**Answer:** C

#### Explanation:

The Snowflake URL type used by directory tables is the scoped URL. This type of URL provides access to files in a stage with metadata, such as the Snowflake file URL, for each file

#### NEW QUESTION 215

- (Topic 3)

What is the difference between a stored procedure and a User-Defined Function (UDF)?

- A. Stored procedures can execute database operations while UDFs cannot.
- B. Returning a value is required in a stored procedure while returning values in a UDF is optional.
- C. Values returned by a stored procedure can be used directly in a SQL statement while the values returned by a UDF cannot.
- D. Multiple stored procedures can be called as part of a single executable statement while a single SQL statement can only call one UDF at a time.

**Answer:** A

#### Explanation:

Stored procedures in Snowflake can perform a variety of database operations, including DDL and DML, whereas UDFs are designed to return values and cannot execute database operations<sup>1</sup>.

#### NEW QUESTION 218

- (Topic 3)

Which of the following activities consume virtual warehouse credits in the Snowflake environment? (Choose two.)

- A. Caching query results
- B. Running EXPLAIN and SHOW commands
- C. Cloning a database
- D. Running a custom query
- E. Running COPY commands

**Answer:** BD

#### Explanation:

Running EXPLAIN and SHOW commands, as well as running a custom query, consume virtual warehouse credits in the Snowflake environment. These activities require computational resources, and therefore, credits are used to account for the usage of these resources. References: [COF-C02] SnowPro Core Certification Exam Study Guide

#### NEW QUESTION 222

- (Topic 3)

Credit charges for Snowflake virtual warehouses are calculated based on which of the following considerations? (Choose two.)

- A. The number of queries executed
- B. The number of active users assigned to the warehouse
- C. The size of the virtual warehouse
- D. The length of time the warehouse is running
- E. The duration of the queries that are executed

**Answer:** CD

#### Explanation:

Credit charges for Snowflake virtual warehouses are calculated based on the size of the virtual warehouse and the length of time the warehouse is running. The size determines the compute resources available, and charges are incurred for the time these resources are utilized

#### NEW QUESTION 226

- (Topic 3)

What happens to the shared objects for users in a consumer account from a share, once a database has been created in that account?

- A. The shared objects are transferred.
- B. The shared objects are copied.
- C. The shared objects become accessible.
- D. The shared objects can be re-shared.

**Answer:** C

#### Explanation:

Once a database has been created in a consumer account from a share, the shared objects become accessible to users in that account. The shared objects are not transferred or copied; they remain in the provider's account and are accessible to the consumer account

#### NEW QUESTION 228

- (Topic 3)

How often are the Account and Table master keys automatically rotated by Snowflake?

- A. 30 Days
- B. 60 Days
- C. 90 Days
- D. 365 Days.

**Answer:** A

#### Explanation:

Snowflake automatically rotates the Account and Table master keys when they are more than 30 days old. Active keys are retired, and new keys are created, ensuring robust security through frequent key changes<sup>1</sup>

#### NEW QUESTION 233

- (Topic 3)

A Snowflake user has two tables that contain numeric values and is trying to find out which values are present in both tables. Which set operator should be used?

- A. INTERSECT
- B. MFRCK
- C. MINUS
- D. UNION

**Answer:** A

#### Explanation:

To find out which numeric values are present in both tables, the INTERSECT set operator should be used. This operator returns rows from one query's result set which also appear in another query's result set, effectively finding the common elements between the two tables<sup>45</sup>.

#### NEW QUESTION 234

- (Topic 3)

What is cached during a query on a virtual warehouse?

- A. All columns in a micro-partition
- B. Any columns accessed during the query
- C. The columns in the result set of the query
- D. All rows accessed during the query

**Answer:** C

#### Explanation:

During a query on a virtual warehouse, the columns in the result set of the query are cached. This allows for faster retrieval of data if the same or a similar query is run again, as the system can retrieve the data from the cache rather than reprocessing the entire query. References: [COF-C02] SnowPro Core Certification Exam Study Guide

#### NEW QUESTION 235

- (Topic 3)

Which Snowflake tool would be BEST to troubleshoot network connectivity?

- A. SnowCLI
- B. SnowUI
- C. SnowSQL
- D. SnowCD

**Answer:** D

#### Explanation:

SnowCD (Snowflake Connectivity Diagnostic Tool) is the best tool provided by Snowflake for troubleshooting network connectivity issues. It helps diagnose and resolve issues related to connecting to Snowflake services [https://docs.snowflake.com/en/user-guide/snowcd.html#:~:text=SnowCD%20\(i.e.%20Snowflake%20Connectivity%20Diagnosti%20c,their%20network%20connection%20to%20Snowflake.](https://docs.snowflake.com/en/user-guide/snowcd.html#:~:text=SnowCD%20(i.e.%20Snowflake%20Connectivity%20Diagnosti%20c,their%20network%20connection%20to%20Snowflake.)

#### NEW QUESTION 237

- (Topic 3)

A Snowflake user has been granted the create data EXCHANGE listing privilege with their role.

Which tasks can this user now perform on the Data Exchange? (Select TWO).

- A. Rename listings.
- B. Delete provider profiles.
- C. Modify listings properties.
- D. Modify incoming listing access requests.
- E. Submit listings for approval/publishing.

**Answer:** CE

#### Explanation:

With the create data EXCHANGE listing privilege, a Snowflake user can modify the properties of listings and submit them for approval or publishing on the Data Exchange. This allows them to manage and share data sets with consumers effectively. References: Based on general data exchange practices in cloud services

as of 2021.

**NEW QUESTION 238**

- (Topic 3)

Which languages require that User-Defined Function (UDF) handlers be written inline? (Select TWO).

- A. Java
- B. Javascript
- C. Scala
- D. Python
- E. SQL

**Answer:** BE

**Explanation:**

User-Defined Function (UDF) handlers must be written inline for Javascript and SQL. These languages allow the UDF logic to be included directly within the SQL statement that creates the UDF.

**NEW QUESTION 243**

- (Topic 4)

What is the purpose of the Snowflake SPLIT TO\_TABLE function?

- A. To count the number of characters in a string
- B. To split a string into an array of sub-strings
- C. To split a string and flatten the results into rows
- D. To split a string and flatten the results into columns

**Answer:** C

**Explanation:**

The purpose of the Snowflake SPLIT\_TO\_TABLE function is to split a string based on a specified delimiter and flatten the results into rows. This table function is useful for transforming a delimited string into a set of rows that can be further processed or queried.

**NEW QUESTION 247**

- (Topic 4)

What tasks can an account administrator perform in the Data Exchange? (Select TWO).

- A. Add and remove members.
- B. Delete data categories.
- C. Approve and deny listing approval requests.
- D. Transfer listing ownership.
- E. Transfer ownership of a provider profile.

**Answer:** AC

**Explanation:**

An account administrator in the Data Exchange can perform tasks such as adding and removing members and approving or denying listing approval requests. These tasks are part of managing the Data Exchange and ensuring that only authorized listings and members are part of it.

**NEW QUESTION 250**

- (Topic 4)

How can a Snowflake administrator determine which user has accessed a database object that contains sensitive information?

- A. Review the granted privileges to the database object.
- B. Review the row access policy for the database object.
- C. Query the ACCESS\_HISTORY view in the ACCOUNT\_USAGE schema.
- D. Query the REPLICATION USAGE HISTORY view in the ORGANIZATION\_USAGE schema.

**Answer:** C

**Explanation:**

To determine which user has accessed a database object containing sensitive information, a Snowflake administrator can query the ACCESS\_HISTORY view in the ACCOUNT\_USAGE schema, which provides information about access to database objects.

**NEW QUESTION 254**

- (Topic 4)

A permanent table and temporary table have the same name, TBL1, in a schema. What will happen if a user executes select \* from TBL1 ;?

- A. The temporary table will take precedence over the permanent table.
- B. The permanent table will take precedence over the temporary table.
- C. An error will say there cannot be two tables with the same name in a schema.
- D. The table that was created most recently will take precedence over the older table.

**Answer:** A

**Explanation:**

In Snowflake, if a temporary table and a permanent table have the same name within the same schema, the temporary table takes precedence over the

permanent table within the session where the temporary table was created4.

#### NEW QUESTION 255

- (Topic 4)

A user wants to access files stored in a stage without authenticating into Snowflake. Which type of URL should be used?

- A. File URL
- B. Staged URL
- C. Scoped URL
- D. Pre-signed URL

**Answer:** D

#### Explanation:

A Pre-signed URL should be used to access files stored in a Snowflake stage without requiring authentication into Snowflake. Pre-signed URLs are simple HTTPS URLs that provide temporary access to a file via a web browser, using a pre-signed access token. The expiration time for the access token is configurable, and this type of URL allows users or applications to directly access or download the files without needing to authenticate into Snowflake5.

References: [COF-C02] SnowPro Core Certification Exam Study Guide

#### NEW QUESTION 260

- (Topic 4)

What will prevent unauthorized access to a Snowflake account from an unknown source?

- A. Network policy
- B. End-to-end encryption
- C. Multi-Factor Authentication (MFA)
- D. Role-Based Access Control (RBAC)

**Answer:** A

#### Explanation:

A network policy in Snowflake is used to restrict access to the Snowflake account from unauthorized or unknown sources. It allows administrators to specify allowed IP address ranges, thus preventing access from any IP addresses not listed in the policy1.

#### NEW QUESTION 264

- (Topic 4)

A user with which privileges can create or manage other users in a Snowflake account? (Select TWO).

- A. GRANT
- B. SELECT
- C. MODIFY
- D. OWNERSHIP
- E. CREATE USER

**Answer:** DE

#### Explanation:

A user with the OWNERSHIP privilege on a user object or the CREATE USER privilege on the account can create or manage other users in a Snowflake account56.

#### NEW QUESTION 267

- (Topic 4)

A tag object has been assigned to a table (TABLE\_A) in a schema within a Snowflake database. Which CREATE object statement will automatically assign the TABLE\_A tag to a target object?

- A. CREATE TABLE <table\_name> LIKE TABLE\_A;
- B. CREATE VIEW <view\_name> AS SELECT \* FROM TABLE\_A;
- C. CREATE TABLE <table\_name> AS SELECT \* FROM TABLE\_A;
- D. CREATE MATERIALIZED VIEW <view name> AS SELECT \* FROM TABLE A;

**Answer:** C

#### Explanation:

When a tag object is assigned to a table, using the statement CREATE TABLE <table\_name> AS SELECT \* FROM TABLE\_A will automatically assign the TABLE\_A tag to the newly created table2.

#### NEW QUESTION 269

- (Topic 4)

Which commands are restricted in owner's rights stored procedures? (Select TWO).

- A. SHOW
- B. MERGE
- C. INSERT
- D. DELETE
- E. DESCRIBE

**Answer:** AE

**Explanation:**

In owner's rights stored procedures, certain commands are restricted to maintain security and integrity. The SHOW and DESCRIBE commands are limited because they can reveal metadata and structure information that may not be intended for all roles.

**NEW QUESTION 270**

- (Topic 4)

Which Snowflake feature allows administrators to identify unused data that may be archived or deleted?

- A. Access history
- B. Data classification
- C. Dynamic Data Masking
- D. Object tagging

**Answer:** A

**Explanation:**

The Access History feature in Snowflake allows administrators to track data access patterns and identify unused data. This information can be used to make decisions about archiving or deleting data to optimize storage and reduce costs.

**NEW QUESTION 271**

- (Topic 4)

How can performance be optimized for a query that returns a small amount of data from a very large base table?

- A. Use clustering keys
- B. Create materialized views
- C. Use the search optimization service
- D. Use the query acceleration service

**Answer:** C

**Explanation:**

The search optimization service in Snowflake is designed to improve the performance of selective point lookup queries on large tables, which is ideal for scenarios where a query returns a small amount of data from a very large base table<sup>1</sup>. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 276**

- (Topic 4)

Which type of loop requires a BREAK statement to stop executing?

- A. FOR
- B. LOOP
- C. REPEAT
- D. WHILE

**Answer:** B

**Explanation:**

The LOOP type of loop in Snowflake Scripting does not have a built-in termination condition and requires a BREAK statement to stop executing<sup>4</sup>.

**NEW QUESTION 278**

- (Topic 4)

What factors impact storage costs in Snowflake? (Select TWO).

- A. The account type
- B. The storage file format
- C. The cloud region used by the account
- D. The type of data being stored
- E. The cloud platform being used

**Answer:** AC

**Explanation:**

The factors that impact storage costs in Snowflake include the account type (Capacity or On Demand) and the cloud region used by the account. These factors determine the rate at which storage is billed, with different regions potentially having different rates<sup>3</sup>.

**NEW QUESTION 283**

- (Topic 4)

What are key characteristics of virtual warehouses in Snowflake? (Select TWO).

- A. Warehouses that are multi-cluster can have nodes of different sizes.
- B. Warehouses can be started and stopped at any time.
- C. Warehouses can be resized at any time, even while running.
- D. Warehouses are billed on a per-minute usage basis.
- E. Warehouses can only be used for querying and cannot be used for data loading.

**Answer:** BC

**Explanation:**

Virtual warehouses in Snowflake can be started and stopped at any time, providing flexibility in managing compute resources. They can also be resized at any time, even while running, to accommodate varying workloads<sup>9</sup><sup>10</sup>. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 286**

- (Topic 4)

What feature of Snowflake Continuous Data Protection can be used for maintenance of historical data?

- A. Access control
- B. Fail-safe
- C. Network policies
- D. Time Travel

**Answer:** D

**Explanation:**

Snowflake's Time Travel feature is used for the maintenance of historical data, allowing users to access and restore data that has been changed or deleted within a defined period<sup>4</sup>.

**NEW QUESTION 291**

- (Topic 4)

Which ACCOUNT\_USAGE schema database role provides visibility into policy-related information?

- A. USAGE\_VIEWER
- B. GOVERNANCE\_VIEWER
- C. OBJECT\_VIEWER
- D. SECURITY\_VIEWER

**Answer:** B

**Explanation:**

The GOVERNANCE\_VIEWER role in the ACCOUNT\_USAGE schema provides visibility into policy-related information within Snowflake. This role is specifically designed to access views that display object metadata and usage metrics related to governance<sup>12</sup>.

**NEW QUESTION 292**

- (Topic 4)

Who can grant object privileges in a regular schema?

- A. Object owner
- B. Schema owner
- C. Database owner
- D. SYSADMIN

**Answer:** A

**Explanation:**

In a regular schema within Snowflake, the object owner has the privilege to grant object privileges. The object owner is typically the role that created the object or to whom the ownership of the object has been transferred<sup>7</sup><sup>8</sup>.

References = [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 296**

- (Topic 4)

Which command is used to start configuring Snowflake for Single Sign-On (SSO)?

- A. CREATE SESSION POLICY
- B. CREATE NETWORK RULE
- C. CREATE SECURITY INTEGRATION
- D. CREATE PASSWORD POLICY

**Answer:** C

**Explanation:**

To start configuring Snowflake for Single Sign-On (SSO), the CREATE SECURITY INTEGRATION command is used. This command sets up a security integration object in Snowflake, which is necessary for enabling SSO with external identity providers using SAML 2.0<sup>1</sup>.

References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 300**

- (Topic 4)

At what level is the MIN\_DATA\_RETENTION\_TIME\_IN\_DAYS parameter set?

- A. Account
- B. Database
- C. Schema
- D. Table

**Answer:** A

**Explanation:**

The MIN\_DATA\_RETENTION\_TIME\_IN\_DAYS parameter is set at the account level. This parameter determines the minimum number of days Snowflake retains historical data for Time Travel operations

### NEW QUESTION 303

- (Topic 4)

Which Snowflake data types can be used to build nested hierarchical data? (Select TWO)

- A. INTEGER
- B. OBJECT
- C. VARIANT
- D. VARCHAR
- E. LIST

**Answer:** BC

#### Explanation:

The Snowflake data types that can be used to build nested hierarchical data are OBJECT and VARIANT. These data types support the storage and querying of semi-structured data, allowing for the creation of complex, nested data structures

### NEW QUESTION 304

- (Topic 4)

While working with unstructured data, which file function generates a Snowflake-hosted file URL to a staged file using the stage name and relative file path as inputs?

- A. GET\_PRESIGNED\_URL
- B. GET\_ABSOLUTE\_PATH
- C. BUILD\_STAGE\_FILE\_URL
- D. BUILD\_SCOPED\_FILE\_URL

**Answer:** C

#### Explanation:

The BUILD\_STAGE\_FILE\_URL function generates a Snowflake-hosted file URL to a staged file using the stage name and relative file path as inputs.

### NEW QUESTION 309

- (Topic 4)

What are the least privileges needed to view and modify resource monitors? (Select TWO).

- A. SELECT
- B. OWNERSHIP
- C. MONITOR
- D. MODIFY
- E. USAGE

**Answer:** CD

#### Explanation:

To view and modify resource monitors, the least privileges needed are MONITOR and MODIFY. These privileges allow a user to monitor credit usage and make changes to resource monitors.

### NEW QUESTION 311

- (Topic 4)

Which Snowflake view is used to support compliance auditing?

- A. ACCESS\_HISTORY
- B. COPY\_HISTORY
- C. QUERY\_HISTORY
- D. ROW\_ACCESS\_POLICIES

**Answer:** A

#### Explanation:

The ACCESS\_HISTORY view in Snowflake is utilized to support compliance auditing. It provides detailed information on data access within Snowflake, including reads and writes by user queries. This view is essential for regulatory compliance auditing as it offers insights into the usage of tables and columns, and maintains a direct link between the user, the query, and the accessed data.

References: [COF-C02] SnowPro Core Certification Exam Study Guide

### NEW QUESTION 314

- (Topic 4)

Which statistics are displayed in a Query Profile that indicate that intermediate results do not fit in memory? (Select TWO).

- A. Bytes scanned
- B. Partitions scanned
- C. Bytes spilled to local storage
- D. Bytes spilled to remote storage
- E. Percentage scanned from cache

**Answer:**

CD

**Explanation:**

The Query Profile statistics that indicate intermediate results do not fit in memory are the bytes spilled to local storage and bytes spilled to remote storage2.

**NEW QUESTION 316**

- (Topic 4)

What is the purpose of the STRIP NULL\_VALUES file format option when loading semi-structured data files into Snowflake?

- A. It removes null values from all columns in the data.
- B. It converts null values to empty strings during loading.
- C. It skips rows with null values during the loading process.
- D. It removes object or array elements containing null values.

**Answer:** D

**Explanation:**

The STRIP NULL\_VALUES file format option, when set to TRUE, removes object or array elements that contain null values during the loading process of semi-structured data files into Snowflake. This ensures that the data loaded into Snowflake tables does not contain these null elements, which can be useful when the ??null?? values in files indicate missing values and have no other special meaning2.

References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 318**

- (Topic 4)

What happens to the objects in a reader account when the DROP MANAGED ACCOUNT command is executed?

- A. The objects are dropped.
- B. The objects enter the Fail-safe period.
- C. The objects enter the Time Travel period.
- D. The objects are immediately moved to the provider account.

**Answer:** A

**Explanation:**

When the DROP MANAGED ACCOUNT command is executed in Snowflake, it removes the managed account, including all objects created within the account, and access to the account is immediately restricted2.

References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 322**

- (Topic 4)

What is the purpose of a Query Profile?

- A. To profile how many times a particular query was executed and analyze its usage statistics over time.
- B. To profile a particular query to understand the mechanics of the query, its behavior, and performance.
- C. To profile the user and/or executing role of a query and all privileges and policies applied on the objects within the query.
- D. To profile which queries are running in each warehouse and identify proper warehouse utilization and sizing for better performance and cost balancing.

**Answer:** B

**Explanation:**

The purpose of a Query Profile is to provide a detailed analysis of a particular query's execution plan, including the mechanics, behavior, and performance. It helps in identifying potential performance bottlenecks and areas for optimization

**NEW QUESTION 323**

- (Topic 4)

Which Snowflake role can manage any object grant globally, including modifying and revoking grants?

- A. USERADMIN
- B. ORGADMIN
- C. SYSADMIN
- D. SECURITYADMIN

**Answer:** D

**Explanation:**

The SECURITYADMIN role in Snowflake can manage any object grant globally, including modifying and revoking grants. This role has the necessary privileges to oversee and control access to all securable objects within the Snowflake environment4.

**NEW QUESTION 324**

- (Topic 5)

How does a Snowflake user extract the URL of a directory table on an external stage for further transformation?

- A. Use the SHOW STAGES command.
- B. Use the DESCRIBE STAGE command.
- C. Use the GET\_ABSOLUTE\_PATH function.
- D. Use the GET\_STAGE\_LOCATION function.

**Answer:** C

**Explanation:**

To extract the URL of a directory table on an external stage for further transformation in Snowflake, the GET\_ABSOLUTE\_PATH function can be used. This function returns the full path of a file or directory within a specified stage, enabling users to dynamically construct URLs for accessing or processing data stored in external stages. References:

? Snowflake Documentation: Working with Stages

**NEW QUESTION 327**

- (Topic 5)

The effects of query pruning can be observed by evaluating which statistics? (Select TWO).

- A. Partitions scanned
- B. Partitions total
- C. Bytes scanned
- D. Bytes read from result
- E. Bytes written

**Answer:** AC

**Explanation:**

Query pruning in Snowflake refers to the optimization technique where the system reduces the amount of data scanned by a query based on the query conditions. This typically involves skipping unnecessary data partitions that do not contribute to the query result. The effectiveness of this technique can be observed through:

? Option A: Partitions scanned. This statistic indicates how many data partitions were actually scanned as a result of query pruning, showing the optimization in action.

? Option C: Bytes scanned. This measures the volume of data physically read during query execution, and a reduction in this number indicates effective query pruning, as fewer bytes are read when unnecessary partitions are skipped.

Options B, D, and E do not directly relate to observing the effects of query pruning. "Partitions total" shows the total available, not the impact of pruning, while "Bytes read from result" and "Bytes written" relate to output rather than the efficiency of data scanning. References: Snowflake documentation on performance tuning and query optimization techniques, specifically how query pruning affects data access.

**NEW QUESTION 332**

- (Topic 5)

Which activities are included in the Cloud Services layer? (Select TWO).

- A. Data storage
- B. Dynamic data masking
- C. Partition scanning
- D. User authentication
- E. Infrastructure management

**Answer:** DE

**Explanation:**

The Cloud Services layer in Snowflake is responsible for a wide range of services that facilitate the management and use of Snowflake, including:

? D. User authentication: This service handles identity and access management, ensuring that only authorized users can access Snowflake resources.

? E. Infrastructure management: This service manages the allocation and scaling of resources to meet user demands, including the management of virtual warehouses, storage, and the orchestration of query execution.

These services are part of Snowflake's fully managed, cloud-based architecture, which abstracts and automates many of the complexities associated with data warehousing. References:

? Snowflake Documentation: Overview of Snowflake Cloud Services

**NEW QUESTION 336**

- (Topic 5)

When referring to User-Defined Function (UDF) names in Snowflake, what does the term overloading mean?

- A. There are multiple SQL UDFs with the same names and the same number of arguments.
- B. There are multiple SQL UDFs with the same names and the same number of argument types.
- C. There are multiple SQL UDFs with the same names but with a different number of arguments or argument types.
- D. There are multiple SQL UDFs with different names but the same number of arguments or argument types.

**Answer:** C

**Explanation:**

In Snowflake, overloading refers to the creation of multiple User-Defined

Functions (UDFs) with the same name but differing in the number or types of their arguments. This feature allows for more flexible function usage, as Snowflake can differentiate between functions based on the context of their invocation, such as the types or the number of arguments passed. Overloading helps to create more adaptable and readable code, as the same function name can be used for similar operations on different types of data.

References:

? Snowflake Documentation: User-Defined Functions

**NEW QUESTION 339**

- (Topic 5)

Which view can be used to determine if a table has frequent row updates or deletes?

- A. TABLES
- B. TABLE\_STORAGE\_METRICS
- C. STORAGE\_DAILY\_HISTORY
- D. STORAGE\_USAGE

**Answer:** B

**Explanation:**

The TABLE\_STORAGE\_METRICS view can be used to determine if a table has frequent row updates or deletes. This view provides detailed metrics on the storage utilization of tables within Snowflake, including metrics that reflect the impact of DML operations such as updates and deletes on table storage. For example, metrics related to the number of active and deleted rows can help identify tables that experience high levels of row modifications, indicating frequent updates or deletions.

References:

? Snowflake Documentation: TABLE\_STORAGE\_METRICS View

**NEW QUESTION 340**

- (Topic 5)

Which function will provide the proxy information needed to protect Snowsight?

- A. SYSTEMADMIN\_TAG
- B. SYSTEM\$GET\_PRIVATELINK
- C. SYSTEMSALLONTLIST
- D. SYSTEMAUTHORIZE

**Answer:** B

**Explanation:**

The SYSTEM\$GET\_PRIVATELINK function in Snowflake provides proxy information necessary for configuring PrivateLink connections, which can protect Snowsight as well as other Snowflake services. PrivateLink enhances security by allowing Snowflake to be accessed via a private connection within a cloud provider's network, reducing exposure to the public internet.

References:

? Snowflake Documentation: PrivateLink Setup

**NEW QUESTION 342**

- (Topic 5)

Which command should be used to unload all the rows from a table into one or more files in a named stage?

- A. COPY INTO
- B. GET
- C. INSERT INTO
- D. PUT

**Answer:** A

**Explanation:**

To unload data from a table into one or more files in a named stage, the COPY INTO <location> command should be used. This command exports the result of a query, such as selecting all rows from a table, into files stored in the specified stage. The COPY INTO command is versatile, supporting various file formats and compression options for efficient data unloading.

References:

? Snowflake Documentation: COPY INTO Location

**NEW QUESTION 345**

- (Topic 5)

If a virtual warehouse runs for 61 seconds, shut down, and then restart and runs for 30 seconds, for how many seconds is it billed?

- A. 60
- B. 91
- C. 120
- D. 121

**Answer:** C

**Explanation:**

Snowflake bills virtual warehouse usage in one-minute increments, rounding up to the nearest minute for any partial minute of compute time used. If a virtual warehouse runs for 61 seconds and then, after being shut down, restarts and runs for an additional 30 seconds, the total time billed would be 120 seconds or 2 minutes. The first 61 seconds are rounded up to 2 minutes, and the subsequent 30 seconds are within a new minute, which is also rounded up to the nearest minute.

References:

? Snowflake Documentation: Virtual Warehouses Billing

**NEW QUESTION 349**

- (Topic 5)

Which statement accurately describes Snowflake's architecture?

- A. It uses a local data repository for all compute nodes in the platform.
- B. It is a blend of shared-disk and shared-everything database architectures.
- C. It is a hybrid of traditional shared-disk and shared-nothing database architectures.
- D. It reorganizes loaded data into internal optimized, compressed, and row-based format.

**Answer:** C

**Explanation:**

Snowflake's architecture is unique in that it combines elements of both traditional shared-disk and shared-nothing database architectures. This hybrid approach

allows Snowflake to offer the scalability and performance benefits of a shared-nothing architecture (with compute and storage separated) while maintaining the simplicity and flexibility of a shared-disk architecture in managing data across all nodes in the system. This results in an architecture that provides on-demand scalability, both vertically and horizontally, without sacrificing performance or data cohesion.

References:

? Snowflake Documentation: Snowflake Architecture

#### NEW QUESTION 352

- (Topic 5)

Which data types optimally store semi-structured data? (Select TWO).

- A. ARRAY
- B. CHARACTER
- C. STRING
- D. VARCHAR
- E. VARIANT

**Answer:** AE

#### Explanation:

In Snowflake, semi-structured data is optimally stored using specific data types that are designed to handle the flexibility and complexity of such data. The VARIANT data type can store structured and semi-structured data types, including JSON, Avro, ORC, Parquet, or XML, in a single column. The ARRAY data type, on the other hand, is suitable for storing ordered sequences of elements, which can be particularly useful for semi-structured data types like JSON arrays. These data types provide the necessary flexibility to store and query semi-structured data efficiently in Snowflake.

References:

? Snowflake Documentation: Semi-structured Data Types

#### NEW QUESTION 354

- (Topic 5)

Which Snowflake data type is used to store JSON key value pairs?

- A. TEXT
- B. BINARY
- C. STRING
- D. VARIANT

**Answer:** D

#### Explanation:

The VARIANT data type in Snowflake is used to store JSON key-value pairs along with other semi-structured data formats like AVRO, BSON, and XML. The VARIANT data type allows for flexible and dynamic data structures within a single column, accommodating complex and nested data. This data type is crucial for handling semi-structured data in Snowflake, enabling users to perform SQL operations on JSON objects and arrays directly.

References:

? Snowflake Documentation: Semi-structured Data Types

#### NEW QUESTION 356

- (Topic 5)

Regardless of which notation is used, what are considerations for writing the column name and element names when traversing semi-structured data?

- A. The column name and element names are both case-sensitive.
- B. The column name and element names are both case-insensitive.
- C. The column name is case-sensitive but element names are case-insensitive.
- D. The column name is case-insensitive but element names are case-sensitive.

**Answer:** D

#### Explanation:

When querying semi-structured data in Snowflake, the behavior towards case sensitivity is distinct between column names and the names of elements within the semi-structured data. Column names follow the general SQL norm of being case-insensitive, meaning you can reference them in any case without affecting the query. However, element names within JSON, XML, or other semi-structured data are case-sensitive. This distinction is crucial for accurate data retrieval and manipulation in Snowflake, especially when working with JSON objects where the case of keys can significantly alter the outcome of queries.

References:

? Snowflake Documentation: Querying Semi-structured Data

#### NEW QUESTION 358

- (Topic 5)

While clustering a table, columns with which data types can be used as clustering keys? (Select TWO).

- A. BINARY
- B. GEOGRAPHY
- C. GEOMETRY
- D. OBJECT
- E. VARIANT

**Answer:** AC

#### Explanation:

A clustering key can be defined when a table is created by appending a CLUSTER Where each clustering key consists of one or more table columns/expressions, which can be of any data type, except GEOGRAPHY, VARIANT, OBJECT, or ARRAY <https://docs.snowflake.com/en/user-guide/tables-clustering-keys>

### NEW QUESTION 363

- (Topic 5)

There are two Snowflake accounts in the same cloud provider region: one is production and the other is non-production. How can data be easily transferred from the production account to the non-production account?

- A. Clone the data from the production account to the non-production account.
- B. Create a data share from the production account to the non-production account.
- C. Create a subscription in the production account and have it publish to the non- production account.
- D. Create a reader account using the production account and link the reader account to the non-production account.

**Answer: B**

#### Explanation:

To easily transfer data from a production account to a non-production account in Snowflake within the same cloud provider region, creating a data share is the most efficient approach. Data sharing allows for live, read-only access to selected data objects from the production account to the non-production account without the need to duplicate or move the actual data. This method facilitates seamless access to the data for development, testing, or analytics purposes in the non-production environment. References:

? Snowflake Documentation: Data Sharing

### NEW QUESTION 364

- (Topic 5)

What does the worksheet and database explorer feature in Snowsight allow users to do?

- A. Add or remove users from a worksheet.
- B. Move a worksheet to a folder or a dashboard.
- C. Combine multiple worksheets into a single worksheet.
- D. Tag frequently accessed worksheets for ease of access.

**Answer: D**

#### Explanation:

The worksheet and database explorer feature in Snowsight allows users to tag frequently accessed worksheets for ease of access. This functionality helps users organize and quickly navigate to the worksheets they use most often, enhancing productivity and streamlining the data exploration and analysis process within Snowsight, Snowflake's web-based query and visualization interface.

References:

? Snowflake Documentation: Snowsight (UI for Snowflake)

### NEW QUESTION 367

- (Topic 5)

Which function is used to convert rows in a relational table to a single VARIANT column?

- A. ARRAY\_AGG
- B. OBJECT\_AGG
- C. ARRAY\_CONSTRUCT
- D. OBJECT\_CONSTRUCT

**Answer: D**

#### Explanation:

The OBJECT\_CONSTRUCT function in Snowflake is used to convert rows in a relational table into a single VARIANT column that represents each row as a JSON object. This function dynamically creates a JSON object from a list of key-value pairs, where each key is a column name and each value is the corresponding column value for a row. This is particularly useful for aggregating and transforming structured data into semi- structured JSON format for further processing or analysis.

References:

? Snowflake Documentation: Semi-structured Data Functions

### NEW QUESTION 368

- (Topic 5)

What does the TableScan operator represent in the Query Profile?

- A. The access to a single table
- B. The access to data stored in stage objects
- C. The list of values provided with the VALUES clause
- D. The records generated using the TABLE (GENERATOR (...)) construct

**Answer: A**

#### Explanation:

In the Query Profile of Snowflake, the TableScan operator represents the access to a single table. This operator indicates that the query execution involved reading data from a table stored in Snowflake. TableScan is a fundamental operation in query execution plans, showing how the database engine retrieves data directly from tables as part of processing a query.

References:

? Snowflake Documentation: Understanding the Query Profile

### NEW QUESTION 370

- (Topic 5)

User1, who has the SYSADMIN role, executed a query on Snowsight. User2, who is in the same Snowflake account, wants to view the result set of the query executed by User1 using the Snowsight query history.

What will happen if User2 tries to access the query history?

- A. If User2 has the sysadmin role they will be able to see the results.
- B. If User2 has the securityadmin role they will be able to see the results.
- C. If User2 has the ACCOUNTADMIN role they will be able to see the results.
- D. User2 will be unable to view the result set of the query executed by User1.

**Answer:** C

**Explanation:**

In Snowflake, the query history and the results of queries executed by a user are accessible based on the roles and permissions. If User1 executed a query with the SYSADMIN role, User2 would be able to view the result set of that query executed by User1 only if User2 has the ACCOUNTADMIN role. The ACCOUNTADMIN role has the broadest set of privileges, including the ability to access all aspects of the account's operation, data, and query history, thus enabling User2 to view the results of queries executed by other users.

References:

? Snowflake Documentation: Understanding Snowflake Roles

**NEW QUESTION 373**

- (Topic 5)

Why would a Snowflake user decide to use a materialized view instead of a regular view?

- A. The base tables do not change frequently.
- B. The results of the view change often.
- C. The query is not resource intensive.
- D. The query results are not used frequently.

**Answer:** A

**Explanation:**

A Snowflake user would decide to use a materialized view instead of a regular view primarily when the base tables do not change frequently. Materialized views store the result of the view query and update it as the underlying data changes, making them ideal for situations where the data is relatively static and query performance is critical. By precomputing and storing the query results, materialized views can significantly reduce query execution times for complex aggregations, joins, and calculations.

References:

? Snowflake Documentation: Materialized Views

**NEW QUESTION 374**

- (Topic 5)

Which service or feature in Snowflake is used to improve the performance of certain types of lookup and analytical queries that use an extensive set of WHERE conditions?

- A. Data classification
- B. Query acceleration service
- C. Search optimization service
- D. Tagging

**Answer:** C

**Explanation:**

The Search Optimization Service in Snowflake is designed to improve the performance of specific types of queries, particularly those involving extensive sets of WHERE conditions. By maintaining a search index on tables, this service can accelerate lookup and analytical queries, making it a valuable feature for optimizing query performance and reducing execution times for complex searches.

References:

? Snowflake Documentation: Search Optimization Service

**NEW QUESTION 376**

- (Topic 5)

What Snowflake database object is derived from a query specification, stored for later use, and can speed up expensive aggregation on large data sets?

- A. Temporary table
- B. External table
- C. Secure view
- D. Materialized view

**Answer:** D

**Explanation:**

A materialized view in Snowflake is a database object derived from a query specification, stored for later use, and can significantly speed up expensive aggregations on large data sets. Materialized views store the result of their underlying query, reducing the need to recompute the result each time the view is accessed. This makes them ideal for improving the performance of read-heavy, aggregate-intensive queries.

References:

? Snowflake Documentation: Using Materialized Views

**NEW QUESTION 379**

- (Topic 5)

Which command removes a role from another role or a user in Snowflake?

- A. ALTER ROLE
- B. REVOKE ROLE
- C. USE ROLE
- D. USE SECONDARY ROLES

**Answer:** B

**Explanation:**

The REVOKE ROLE command is used to remove a role from another role or a user in Snowflake. This command is part of Snowflake's role-based access control system, allowing administrators to manage permissions and access to database objects efficiently by adding or removing roles from users or other roles.

References:

? Snowflake Documentation: REVOKE ROLE

**NEW QUESTION 383**

- (Topic 5)

What criteria does Snowflake use to determine the current role when initiating a session? (Select TWO).

- A. If a role was specified as part of the connection and that role has been granted to the Snowflake user, the specified role becomes the current role.
- B. If no role was specified as part of the connection and a default role has been defined for the Snowflake user, that role becomes the current role.
- C. If no role was specified as part of the connection and a default role has not been set for the Snowflake user, the session will not be initiated and the log in will fail.
- D. If a role was specified as part of the connection and that role has not been granted to the Snowflake user, it will be ignored and the default role will become the current role.
- E. If a role was specified as part of the connection and that role has not been granted to the Snowflake user, the role is automatically granted and it becomes the current role.

**Answer:** AB

**Explanation:**

When initiating a session in Snowflake, the system determines the current role based on the user's connection details and role assignments. If a user specifies a role during the connection, and that role is already granted to them, Snowflake sets it as the current role for the session. Alternatively, if no role is specified during the connection, but the user has a default role assigned, Snowflake will use this default role as the current session role. These mechanisms ensure that users operate within their permissions, enhancing security and governance within Snowflake environments.

References:

? Snowflake Documentation: Understanding Roles

**NEW QUESTION 384**

- (Topic 5)

Which Snowflake mechanism is used to limit the number of micro-partitions scanned by a query?

- A. Caching
- B. Cluster depth
- C. Query pruning
- D. Retrieval optimization

**Answer:** C

**Explanation:**

Query pruning in Snowflake is the mechanism used to limit the number of micro-partitions scanned by a query. By analyzing the filters and conditions applied in a query, Snowflake can skip over micro-partitions that do not contain relevant data, thereby reducing the amount of data processed and improving query performance. This technique is particularly effective for large datasets and is a key component of Snowflake's performance optimization features.

References:

? Snowflake Documentation: Query Performance Optimization

**NEW QUESTION 387**

- (Topic 5)

What is the only supported character set for loading and unloading data from all supported file formats?

- A. UTF-8
- B. UTF-16
- C. ISO-8859-1
- D. WINDOWS-1253

**Answer:** A

**Explanation:**

UTF-8 is the only supported character set for loading and unloading data from all supported file formats in Snowflake. UTF-8 is a widely used encoding that supports a large range of characters from various languages, making it suitable for internationalization and ensuring data compatibility across different systems and platforms.

References:

? Snowflake Documentation: Data Loading and Unloading

**NEW QUESTION 392**

- (Topic 6)

Which Snowflake objects can be restored using Time Travel? (Select VNO).

- A. Roles
- B. Users
- C. Databases
- D. Schemas
- E. Virtual warehouses

**Answer:** CD

**Explanation:**

Snowflake's Time Travel feature allows users to access historical data within a specific period. This feature supports the restoration of various objects, including databases and schemas, to their previous states. Time Travel can be used for recovering dropped objects, undoing accidental changes, or analyzing data changes over time. However, it does not support user or role objects like Users and Roles, or compute resources like Virtual Warehouses. References: Snowflake Documentation on Time Travel

**NEW QUESTION 393**

- (Topic 6)

Which object type is granted permissions for reading a table?

- A. User
- B. Role
- C. Attribute
- D. Schema

**Answer: B**

**Explanation:**

In Snowflake, permissions for accessing database objects, including tables, are not granted directly to users but rather to roles. A role encapsulates a collection of privileges on various Snowflake objects. Users are then granted roles, and through those roles, they inherit the permissions necessary to read a table or perform other actions. This approach adheres to the principle of least privilege, allowing for granular control over database access and simplifying the management of user permissions.

Reference: Snowflake's official documentation on access control introduces the concept of roles and how they are used to manage permissions: <https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#roles>

**NEW QUESTION 394**

- (Topic 6)

Which type of workload traditionally benefits from the use of the query acceleration service?

- A. Workloads with a predictable data volume for each query
- B. Workloads that include on-demand data analyses
- C. Queries with small scans and non-selective filters
- D. Queries that do not have filters or aggregation

**Answer: B**

**Explanation:**

The query acceleration service in Snowflake is beneficial for workloads that include on-demand data analyses. This service optimizes query performance by dynamically allocating additional resources to execute queries faster, particularly useful for ad-hoc analysis where data volume and complexity can vary.

References:

? Snowflake Documentation: Query Acceleration Service

**NEW QUESTION 396**

- (Topic 6)

What do temporary and transient tables have in common in Snowflake? (Select TWO).

- A. Both tables have no Fail-safe period.
- B. Both tables have data retention period maximums of one day.
- C. Both tables are visible only to a single user session.
- D. For both tables the retention period ends when the tables are dropped.
- E. For both tables, the retention period does not end when the session ends

**Answer: AD**

**Explanation:**

Temporary and transient tables in Snowflake share several characteristics, notably, neither table type has a Fail-safe period. Fail-safe is a feature that provides additional data protection beyond the Time Travel period. However, this feature does not apply to temporary or transient tables. Additionally, for both types of tables, the data retention period effectively ends when the tables are dropped. This means that once these tables are deleted, their data is not recoverable, distinguishing them from permanent tables, which benefit from Snowflake's Time Travel and Fail-safe features. References: Snowflake Documentation on Table Types

**NEW QUESTION 398**

- (Topic 6)

Which security models are used in Snowflake to manage access control? (Select TWO).

- A. Discretionary Access Control (DAC)
- B. Identity Access Management (IAM)
- C. Mandatory Access Control (MAC)
- D. Role-Based Access Control (RBAC)
- E. Security Assertion Markup Language (SAML)

**Answer: AD**

**Explanation:**

Snowflake uses both Discretionary Access Control (DAC) and Role-Based Access Control (RBAC) to manage access control. DAC allows object owners to grant access privileges to other users. RBAC assigns permissions to roles, and roles are then granted to users, making it easier to manage permissions based on user roles within the organization.

References:

? Snowflake Documentation: Access Control in Snowflake

#### NEW QUESTION 399

- (Topic 6)

What does Snowflake recommend a user do if they need to connect to Snowflake with a tool or technology that is not listed in Snowflake partner ecosystem?

- A. Use Snowflake's native API.
- B. Use a custom-built connector.
- C. Contact Snowflake Support for a new driver.
- D. Connect through Snowflake's JDBC or ODBC drivers

**Answer: D**

#### Explanation:

If a user needs to connect to Snowflake with a tool or technology that is not listed in Snowflake's partner ecosystem, Snowflake recommends using its JDBC or ODBC drivers. These drivers provide a standard method of connecting from various tools and programming languages to Snowflake, offering wide compatibility and flexibility. By using these drivers, users can establish connections to Snowflake from their applications, ensuring they can leverage the capabilities of Snowflake regardless of the specific tools or technologies they are using. References: Snowflake Documentation on Client Drivers

#### NEW QUESTION 404

- (Topic 6)

Which key access control concept does Snowflake describe as a defined level of access to an object?

- A. Grant
- B. Privilege
- C. Role
- D. Session

**Answer: B**

#### Explanation:

In Snowflake, the term "privilege" refers to a defined level of access to an object. Privileges are specific actions that roles can perform on securable objects in Snowflake, such as tables, views, warehouses, databases, and schemas. These privileges are granted to roles and can be further granted to users through their roles, forming the basis of Snowflake's access control framework. References: Snowflake Documentation on Access Control Privileges

#### NEW QUESTION 408

- (Topic 6)

How does Snowflake utilize clustering information to improve query performance?

- A. It prunes unnecessary micro-partitions based on clustering metadata.
- B. It compresses the data within micro-partitions for faster querying.
- C. It automatically allocates additional resources to improve query execution.
- D. It organizes clustering information to speed-up data retrieval from storage

**Answer: A**

#### Explanation:

Snowflake utilizes clustering information to enhance query performance by pruning unnecessary micro-partitions.

? Clustering Metadata: Snowflake stores clustering information for each micro-partition, which includes data range and distribution.

? Pruning Micro-partitions: When a query is executed, Snowflake uses this clustering metadata to identify and eliminate micro-partitions that do not match the query criteria, thereby reducing the amount of data scanned and improving query performance.

References:

? Snowflake Documentation on Clustering

? Snowflake Documentation on Micro-partition Pruning

#### NEW QUESTION 409

- (Topic 6)

What are characteristics of Snowflake network policies? (Select TWO).

- A. They can be set for any Snowflake Edition.
- B. They can be applied to roles.
- C. They restrict or enable access to specific IP addresses.
- D. They are activated using ALTER DATABASE SQL commands.
- E. They can only be managed using the ORGADMIN role.

**Answer: AC**

#### Explanation:

Snowflake network policies are a security feature that allows administrators to control access to Snowflake by specifying allowed and blocked IP address ranges. These policies apply to all editions of Snowflake, making them widely applicable across different Snowflake environments. They are specifically designed to restrict or enable access based on the originating IP addresses of client requests, adding an extra layer of security. Network policies are not applied to roles but are set at the account or user level. They are not activated using ALTER DATABASE SQL commands but are managed through ALTER ACCOUNT or ALTER NETWORK POLICY commands. The management of network policies does not exclusively require the ORGADMIN role; instead, they can be managed by users with the necessary privileges on the account.

Reference: Snowflake Documentation on Network Policies:

<https://docs.snowflake.com/en/user-guide/network-policies.html>

#### NEW QUESTION 411

- (Topic 6)

Which service or tool is a Command Line Interface (CLI) client used for connecting to Snowflake to execute SQL queries?

- A. Snowsight
- B. SnowCD
- C. Snowpark
- D. SnowSQL

**Answer:** D

**Explanation:**

SnowSQL is the Command Line Interface (CLI) client provided by Snowflake for executing SQL queries and performing various tasks. It allows users to connect to their Snowflake accounts and interact with the Snowflake data warehouse.

? Installation: SnowSQL can be downloaded and installed on various operating systems.

? Configuration: Users need to configure SnowSQL with their Snowflake account credentials.

? Usage: Once configured, users can run SQL queries, manage data, and perform administrative tasks through the CLI.

References:

? Snowflake Documentation: SnowSQL

? Snowflake Documentation: Installing SnowSQL

**NEW QUESTION 416**

- (Topic 6)

Who can activate a network policy for users in a Snowflake account? (Select TWO)

- A. ACCOUNTADMIN
- B. USERADMIN
- C. PUBLIC
- D. SYSADMIN
- E. Any role that has the global ATTACH POLICY privilege

**Answer:** AE

**Explanation:**

Network policies in Snowflake are used to control access to Snowflake accounts based on IP address ranges. These policies can be activated by specific roles that have the necessary privileges.

? Role: ACCOUNTADMIN:

? Role with Global ATTACH POLICY Privilege:

References:

? Snowflake Documentation: Network Policies

**NEW QUESTION 418**

- (Topic 6)

A user needs to MINIMIZE the cost of large tables that are used to store transitory data. The data does not need to be protected against failures, because the data can be reconstructed outside of Snowflake.

What table type should be used?

- A. Permanent
- B. Transient
- C. Temporary
- D. Externa

**Answer:** B

**Explanation:**

For minimizing the cost of large tables that are used to store transitory data, which does not need to be protected against failures because it can be reconstructed outside of Snowflake, the best table type to use is Transient. Transient tables in Snowflake are designed for temporary or transitory data storage and offer reduced storage costs compared to permanent tables. However, unlike temporary tables, they persist across sessions until explicitly dropped.

? Why Transient Tables: Transient tables provide a cost-effective solution for storing data that is temporary but needs to be available longer than a single session. They have lower data storage costs because Snowflake does not maintain historical data (Time Travel) for as long as it does for permanent tables.

? Creating a Transient Table:

```
CREATE TRANSIENT TABLE my_transient_table (...);
```

? Use Case Considerations: Transient tables are ideal for scenarios where the data is not critical, can be easily recreated, and where cost optimization is a priority. They are suitable for development, testing, or staging environments where data longevity is not a concern.

Reference: For more details on transient tables and their usage scenarios, refer to the Snowflake documentation on table types: <https://docs.snowflake.com/en/sql-reference/sql/create-table.html#table-types>

**NEW QUESTION 420**

- (Topic 6)

Use of which file function allows a user to share unstructured data from an internal stage with an external reporting tool that does not have access to Snowflake">

- A. BUILD\_SCOPED\_FILE\_URL
- B. GET\_PRE\_SIGNED\_URL
- C. BUILD\_STAGE\_FILE\_URL
- D. GET\_STAGE\_LOCATION

**Answer:** B

**Explanation:**

The GET\_PRE\_SIGNED\_URL function in Snowflake generates a pre-signed URL for a file in an internal stage. This URL can be shared with external tools or users who do not have direct access to Snowflake, allowing them to download the file.

? Generate Pre-Signed URL:

```
SELECT GET_PRE_SIGNED_URL(@my_stage/file.txt);
```

? Share the URL: The generated URL can be shared with external users or applications, enabling them to access the file directly.

References:

- ? Snowflake Documentation: GET\_PRESIGNED\_URL
- ? Snowflake Documentation: Working with Stages

**NEW QUESTION 421**

- (Topic 6)

Which governance feature is supported by all Snowflake editions?

- A. Object tags
- B. Masking policies
- C. Row access policies
- D. OBJECT\_DEPENDENCIES View

**Answer:** D

**Explanation:**

Snowflake's governance features vary across different editions, but the OBJECT\_DEPENDENCIESview is supported by all Snowflake editions. This feature is part of Snowflake's Information Schema and is designed to help users understand the dependencies between various objects in their Snowflake environment.

TheOBJECT\_DEPENDENCIESview provides a way to query and analyze the relationships and dependencies among different database objects, such as tables, views, and stored procedures. This is crucial for governance, as it allows administrators and data engineers to assess the impact of changes, understand object relationships, and ensure proper management of data assets.

Object tags, masking policies, and row access policies are more advanced features that offer fine-grained data governance capabilities such as tagging objects for classification, dynamically masking sensitive data based on user roles, and controlling row-level access to data. These features may have varying levels of support across different Snowflake editions, with some features being exclusive to higher-tier editions.

Reference: Snowflake Documentation on Information Schema

([https://docs.snowflake.com/en/sql-reference/info-schema/object\\_dependencies.html](https://docs.snowflake.com/en/sql-reference/info-schema/object_dependencies.html))

**NEW QUESTION 422**

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### **COF-C02 Practice Exam Features:**

- \* COF-C02 Questions and Answers Updated Frequently
- \* COF-C02 Practice Questions Verified by Expert Senior Certified Staff
- \* COF-C02 Most Realistic Questions that Guarantee you a Pass on Your First Try
- \* COF-C02 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The COF-C02 Practice Test Here](#)**