



# **Snowflake**

## **Exam Questions COF-C02**

SnowPro Core Certification Exam (COF-C02)

## About ExamBible

*[Your Partner of IT Exam](#)*

## Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

## Our Advances

### \* 99.9% Uptime

All examinations will be up to date.

### \* 24/7 Quality Support

We will provide service round the clock.

### \* 100% Pass Rate

Our guarantee that you will pass the exam.

### \* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

### 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)

The Information Schema and Account Usage Share provide storage information for which of the following objects? (Choose three.)

- A. Users
- B. Tables
- C. Databases
- D. Internal Stages

**Answer:** BCD

#### Explanation:

The Information Schema and Account Usage Share in Snowflake provide metadata and historical usage data for various objects within a Snowflake account. Specifically, they offer storage information for Tables, Databases, and Internal Stages. These schemas contain views and table functions that allow users to query object metadata and usage metrics, such as the amount of data stored and historical activity.

? Tables: The storage information includes data on the daily average amount of data in database tables.

? Databases: For databases, the storage usage is calculated based on all the data contained within the database, including tables and stages.

? Internal Stages: Internal stages are locations within Snowflake for temporarily storing data, and their storage usage is also tracked.

References: The information is verified according to the SnowPro Core Certification Study Guide and Snowflake documentation

### NEW QUESTION 3

- (Topic 1)

What is the default character set used when loading CSV files into Snowflake?

- A. UTF-8
- B. UTF-16
- C. ISO S859-1
- D. ANSI\_X3.A

**Answer:** A

#### Explanation:

[https://docs.snowflake.com/en/user-guide/intro-summary-loading.html#:~:text=For%20delimited%20files%20\(CSV%2C%20TSV,encoding%20to%20use%20for%20loading.](https://docs.snowflake.com/en/user-guide/intro-summary-loading.html#:~:text=For%20delimited%20files%20(CSV%2C%20TSV,encoding%20to%20use%20for%20loading.)

For delimited files (CSV, TSV, etc.), the default character set is UTF-8. To use any other characters sets, you must explicitly specify the encoding to use for loading. For the list of supported character sets, see Supported Character Sets for Delimited Files (in this topic).

### NEW QUESTION 4

- (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 5

- (Topic 1)

What is a limitation of a Materialized View?

- A. A Materialized View cannot support any aggregate functions
- B. A Materialized View can only reference up to two tables
- C. A Materialized View cannot be joined with other tables
- D. A Materialized View cannot be defined with a JOIN

**Answer:** D

**Explanation:**

Materialized Views in Snowflake are designed to store the result of a query and can be refreshed to maintain up-to-date data. However, they have certain limitations, one of which is that they cannot be defined using a JOIN clause. This means that a Materialized View can only be created based on a single source table and cannot combine data from multiple tables using JOIN operations.

References:

- ? Snowflake Documentation on Materialized Views
- ? SnowPro® Core Certification Study Guide

**NEW QUESTION 6**

- (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 Documentation<sup>1</sup>
- ? Data Storage Considerations | Snowflake Documentation<sup>2</sup>
- ? Snowflake SnowPro Core Certification Study Guide<sup>3</sup>
- ? Snowflake Data Cloud Glossary
- <https://docs.snowflake.com/en/user-guide/data-availability.html>

**NEW QUESTION 7**

- (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 8**

- (Topic 1)

Which stage type can be altered and dropped?

- A. Database stage
- B. External stage
- C. Table stage
- D. User stage

**Answer:** B

**Explanation:**

External stages can be altered and dropped in Snowflake. An external stage points to an external location, such as an S3 bucket, where data files are stored.

Users can modify the stage's definition or drop it entirely if it's no longer needed. This is in contrast to table stages, which are tied to specific tables and cannot be altered or dropped independently.

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Snowflake Documentation on Stages<sup>1</sup>

#### NEW QUESTION 9

- (Topic 1)

Which of the following describes external functions in Snowflake?

- A. They are a type of User-defined Function (UDF).
- B. They contain their own SQL code.
- C. They call code that is stored inside of Snowflake.
- D. They can return multiple rows for each row received

**Answer:** A

#### Explanation:

External functions in Snowflake are a special type of User-Defined Function (UDF) that call code executed outside of Snowflake, typically through a remote service. Unlike traditional UDFs, external functions do not contain SQL code within Snowflake; instead, they interact with external services to process data.  
<https://docs.snowflake.com/en/sql-reference/external-functions.html#:~:text=External%20functions%20are%20user%2Ddefined,code%20running%20outside%20of%20Snowflake.>

#### NEW QUESTION 10

- (Topic 1)

True or False: Reader Accounts are able to extract data from shared data objects for use outside of Snowflake.

- A. True
- B. False

**Answer:** B

#### Explanation:

Reader accounts in Snowflake are designed to allow users to read data shared with them but do not have the capability to extract data for use outside of Snowflake. They are intended for consuming shared data within the Snowflake environment only.

#### 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 Duo.

#### NEW QUESTION 13

- (Topic 1)

True or False: A Virtual Warehouse can be resized while suspended.

- A. True
- B. False

**Answer:** A

#### Explanation:

Virtual Warehouses in Snowflake can indeed be resized while they are suspended. Resizing a warehouse involves changing the number of compute resources (servers) allocated to it, which can be done to adjust performance and cost. When a warehouse is suspended, it is not currently running any queries, but its definition and metadata remain intact, allowing for modifications like resizing.

Reference: <https://docs.snowflake.com/en/user-guide/warehouses-tasks.html#effects-of-resizing-a-suspended-warehouse>

#### NEW QUESTION 17

- (Topic 1)

What are value types that a VARIANT column can store? (Select TWO)

- A. STRUCT
- B. OBJECT
- C. BINARY
- D. ARRAY
- E. CLOB

**Answer:** BD

#### Explanation:

A VARIANT column in Snowflake can store semi-structured data types. This includes:

? B. OBJECT: An object is a collection of key-value pairs in JSON, and a VARIANT column can store this type of data structure.

? D. ARRAY: An array is an ordered list of zero or more values, which can be of any variant-supported data type, including objects or other arrays.

The VARIANT data type is specifically designed to handle semi-structured data like JSON, Avro, ORC, Parquet, or XML, allowing for the storage of nested and complex data structures.

References:

? Snowflake Documentation on Semi-Structured Data Types

? SnowPro® Core Certification Study Guide

#### NEW QUESTION 20

- (Topic 1)

Which services does the Snowflake Cloud Services layer manage? (Select TWO).

A. Compute resources

B. Query execution

C. Authentication

D. Data storage

E. Metadata

**Answer:** CE

#### Explanation:

The Snowflake Cloud Services layer manages a variety of services that are crucial for the operation of the Snowflake platform. Among these services, Authentication and Metadata management are key components. Authentication is essential for controlling access to the Snowflake environment, ensuring that only authorized users can perform actions within the platform. Metadata management involves handling all the metadata related to objects within Snowflake, such as tables, views, and databases, which is vital for the organization and retrieval of data.

References:

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

? Snowflake Documentation12 <https://docs.snowflake.com/en/user-guide/intro-key-concepts.html>

#### NEW QUESTION 23

- (Topic 1)

Which of the following describes how multiple Snowflake accounts in a single organization relate to various cloud providers?

A. Each Snowflake account can be hosted in a different cloud vendor and region.

B. Each Snowflake account must be hosted in a different cloud vendor and region

C. All Snowflake accounts must be hosted in the same cloud vendor and region

D. Each Snowflake account can be hosted in a different cloud vendor, but must be in the same region.

**Answer:** A

#### Explanation:

Snowflake's architecture allows for flexibility in account hosting across different cloud vendors and regions. This means that within a single organization, different Snowflake accounts can be set up in various cloud environments, such as AWS, Azure, or GCP, and in different geographical regions. This allows organizations to leverage the global infrastructure of multiple cloud providers and optimize their data storage and computing needs based on regional requirements, data sovereignty laws, and other considerations.

<https://docs.snowflake.com/en/user-guide/intro-regions.html>

#### NEW QUESTION 26

- (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 28

- (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 30

- (Topic 1)

Where would a Snowflake user find information about query activity from 90 days ago?

- A. account usage . query history view
- B. account usage.query history archive View
- C. information schema . cruery\_history view
- D. information schema - query history\_by\_session view

**Answer:** B

#### Explanation:

To find information about query activity from 90 days ago, a Snowflake user should use the account\_usage.query\_history\_archive view. This view is designed to provide access to historical query data beyond the default 14-day retention period found in the standard query\_history view. It allows users to analyze and audit past query activities for up to 365 days after the date of execution, which includes the 90-day period mentioned. References:

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

? Snowflake Documentation on Account Usage Schema1

### NEW QUESTION 31

- (Topic 1)

What are the default Time Travel and Fail-safe retention periods for transient tables?

- A. Time Travel - 1 da
- B. Fail-safe - 1 day
- C. Time Travel - 0 day
- D. Fail-safe - 1 day
- E. Time Travel - 1 da
- F. Fail-safe - 0 days
- G. Transient tables are retained in neither Fail-safe nor Time Travel

**Answer:** C

#### Explanation:

Transient tables in Snowflake have a default Time Travel retention period of 1 day, which allows users to access historical data within the last 24 hours. However, transient tables do not have a Fail-safe period. Fail-safe is an additional layer of data protection that retains data beyond the Time Travel period for recovery purposes in case of extreme data loss. Since transient tables are designed for temporary or intermediate workloads with no requirement for long-term durability, they do not include a Fail-safe period by default1.

References:

? Snowflake Documentation on Storage Costs for Time Travel and Fail-safe

### NEW QUESTION 33

- (Topic 1)

Which command can be used to load data into an internal stage?

- A. LOAD
- B. copy
- C. GET
- D. PUT

**Answer:** D

#### Explanation:

The PUT command is used to load data into an internal stage in Snowflake. This command uploads data files from a local file system to a named internal stage, making the data available for subsequent loading into a Snowflake table using the COPY INTO command.

References:

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

? Snowflake Documentation on Data Loading

### NEW QUESTION 38

- (Topic 1)

Which of the following is a valid source for an external stage when the Snowflake account is located on Microsoft Azure?

- A. An FTP server with TLS encryption
- B. An HTTPS server with WebDAV
- C. A Google Cloud storage bucket
- D. A Windows server file share on Azure

**Answer:** D

#### Explanation:

In Snowflake, when the account is located on Microsoft Azure, a valid source for an external stage can be an Azure container or a folder path within an Azure container. This includes Azure Blob storage which is accessible via the azure:// endpoint. A Windows server file share on Azure, if configured properly, can be a valid source for staging data files for Snowflake. Options A, B, and C are not supported as direct sources for an external stage in Snowflake on Azure12.

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

### NEW QUESTION 39

- (Topic 1)

True or False: A 4X-Large Warehouse may, at times, take longer to provision than a X- Small Warehouse.

- A. True
- B. False

**Answer:** A

**Explanation:**

Provisioning time can vary based on the size of the warehouse. A 4X-Large Warehouse typically has more resources and may take longer to provision compared to a X-Small Warehouse, which has fewer resources and can generally be provisioned more quickly. References: Understanding and viewing Fail-safe | Snowflake Documentation

**NEW QUESTION 40**

- (Topic 1)

Which feature is only available in the Enterprise or higher editions of Snowflake?

- A. Column-level security
- B. SOC 2 type II certification
- C. Multi-factor Authentication (MFA)
- D. Object-level access control

**Answer:** A

**Explanation:**

Column-level security is a feature that allows fine-grained control over access to specific columns within a table. This is particularly useful for managing sensitive data and ensuring that only authorized users can view or manipulate certain pieces of information. According to my last update, this feature was available in the Enterprise Edition or higher editions of Snowflake.

References: Based on my internal data as of 2021, column-level security is an advanced feature typically reserved for higher-tiered editions like the Enterprise Edition in data warehousing solutions such as Snowflake.

<https://docs.snowflake.com/en/user-guide/intro-editions.html>

**NEW QUESTION 42**

- (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 46**

- (Topic 1)

True or False: Fail-safe can be disabled within a Snowflake account.

- A. True
- B. False

**Answer:** B

**Explanation:**

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

Separate and distinct from Time Travel, Fail-safe ensures historical data is protected in the event of a system failure or other catastrophic event, e.g. a hardware failure or security breach. Fail-safe feature cannot be enabled or disabled from the user end.

**NEW QUESTION 47**

- (Topic 1)

Which data type can be used to store geospatial data in Snowflake?

- A. Variant
- B. Object
- C. Geometry
- D. Geography

**Answer:** D

**Explanation:**

Snowflake supports two geospatial data

types: GEOGRAPHY and GEOMETRY. The GEOGRAPHY data type is used to store geospatial data that models the Earth as a perfect sphere, which is suitable for global geospatial data. This data type follows the WGS 84 standard and is used for storing points, lines, and polygons on the Earth's surface. The GEOMETRY data type, on the other hand, represents features in a planar (Euclidean, Cartesian) coordinate system and is typically used for local spatial reference systems. Since the question specifically asks about geospatial data, which commonly refers to Earth-related spatial data, the correct answer is GEOGRAPHY.

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

#### NEW QUESTION 51

- (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 Charges<sup>1</sup>

#### NEW QUESTION 53

- (Topic 1)

When reviewing a query profile, what is a symptom that a query is too large to fit into the memory?

- A. A single join node uses more than 50% of the query time
- B. Partitions scanned is equal to partitions total
- C. An AggregateOperator node is present
- D. The query is spilling to remote storage

**Answer: D**

#### Explanation:

When a query in Snowflake is too large to fit into the available memory, it will start spilling to remote storage. This is an indication that the memory allocated for the query is insufficient for its execution, and as a result, Snowflake uses remote disk storage to handle the overflow. This spill to remote storage can lead to slower query performance due to the additional I/O operations required.

References:

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

? Snowflake Documentation on Query Profile<sup>1</sup>

? Snowpro Core Certification Exam Flashcards<sup>2</sup>

#### NEW QUESTION 57

- (Topic 1)

What happens when a virtual warehouse is resized?

- A. When increasing the size of an active warehouse the compute resource for all running and queued queries on the warehouse are affected
- B. When reducing the size of a warehouse the compute resources are removed only when they are no longer being used to execute any current statements.
- C. The warehouse will be suspended while the new compute resource is provisioned and will resume automatically once provisioning is complete.
- D. Users who are trying to use the warehouse will receive an error message until the resizing is complete

**Answer: A**

#### Explanation:

When a virtual warehouse in Snowflake is resized, specifically when it is increased in size, the additional compute resources become immediately available to all running and queued queries. This means that the performance of these queries can improve due to the increased resources. Conversely, when the size of a warehouse is reduced, the compute resources are not removed until they are no longer being used by any current operations<sup>1</sup>.

References:

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

? Snowflake Documentation on Virtual Warehouses<sup>2</sup>

#### NEW QUESTION 58

- (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

### NEW QUESTION 60

- (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 62

- (Topic 1)

Which statement about billing applies to Snowflake credits?

- A. Credits are billed per-minute with a 60-minute minimum
- B. Credits are used to pay for cloud data storage usage
- C. Credits are consumed based on the number of credits billed for each hour that a warehouse runs
- D. Credits are consumed based on the warehouse size and the time the warehouse is running

**Answer:** D

#### Explanation:

Snowflake credits are the unit of measure for the compute resources used in Snowflake. The number of credits consumed depends on the size of the virtual warehouse and the time it is running. Larger warehouses consume more credits per hour than smaller ones, and credits are billed for the time the warehouse is active, regardless of the actual usage within that time.

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

### NEW QUESTION 67

- (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 72

- (Topic 1)

What is the default File Format used in the COPY command if one is not specified?

- A. CSV
- B. JSON
- C. Parquet
- D. XML

**Answer:** A

#### Explanation:

The default file format for the COPY command in Snowflake, when not specified, is CSV (Comma-Separated Values). This format is widely used for data exchange because it is simple, easy to read, and supported by many data analysis tools.

### NEW QUESTION 75

- (Topic 1)

What is the minimum Snowflake edition required to create a materialized view?

- A. Standard Edition
- B. Enterprise Edition

- C. Business Critical Edition
- D. Virtual Private Snowflake Edition

**Answer:** B

**Explanation:**

Materialized views in Snowflake are a feature that allows for the pre-computation and storage of query results for faster query performance. This feature is available starting from the Enterprise Edition of Snowflake. It is not available in the Standard Edition, and while it is also available in higher editions like Business Critical and Virtual Private Snowflake, the Enterprise Edition is the minimum requirement. References:

? Snowflake Documentation on CREATE MATERIALIZED VIEW<sup>1</sup>.

? Snowflake Documentation on Working with Materialized Views <https://docs.snowflake.com/en/sql-reference/sql/create-materialized-view.html#:~:text=Materialized%20views%20require%20Enterprise%20Edition,upgrading%2C%20please%20contact%20Snowflake%20Support>.

**NEW QUESTION 79**

- (Topic 1)

Which of the following describes how clustering keys work in Snowflake?

- A. Clustering keys update the micro-partitions in place with a full sort, and impact the DML operations.
- B. Clustering keys sort the designated columns over time, without blocking DML operations
- C. Clustering keys create a distributed, parallel data structure of pointers to a table's rows and columns
- D. Clustering keys establish a hashed key on each node of a virtual warehouse to optimize joins at run-time

**Answer:** B

**Explanation:**

Clustering keys in Snowflake work by sorting the designated columns over time. This process is done in the background and does not block data manipulation language (DML) operations, allowing for normal database operations to continue without interruption. The purpose of clustering keys is to organize the data within micro-partitions to optimize query performance<sup>1</sup>.

References:

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

? Snowflake Documentation on Clustering<sup>1</sup>

**NEW QUESTION 83**

- (Topic 1)

A company's security audit requires generating a report listing all Snowflake logins (e.g., date and user) within the last 90 days. Which of the following statements will return the required information?

- A. SELECT LAST\_SUCCESS\_LOGIN, LOGIN\_NAME FROM ACCOUNT\_USAGE.USERS;
- B. SELECT EVENT\_TIMESTAMP, USER\_NAME FROM table(information\_schema.login\_history\_by\_user())
- C. SELECT EVENT\_TIMESTAMP, USER\_NAME FROM ACCOUNT\_USAGE.ACCESS\_HISTORY;
- D. SELECT EVENT\_TIMESTAMP, USER\_NAME FROM ACCOUNT\_USAGE.LOGIN\_HISTORY;

**Answer:** D

**Explanation:**

To generate a report listing all Snowflake logins within the last 90 days, the ACCOUNT\_USAGE.LOGIN\_HISTORY view should be used. This view provides information about login attempts, including successful and unsuccessful logins, and is suitable for security audits<sup>4</sup>.

**NEW QUESTION 85**

- (Topic 1)

Which of the following are best practice recommendations that should be considered when loading data into Snowflake? (Select TWO).

- A. Load files that are approximately 25 MB or smaller.
- B. Remove all dates and timestamps.
- C. Load files that are approximately 100-250 MB (or larger)
- D. Avoid using embedded characters such as commas for numeric data types
- E. Remove semi-structured data types

**Answer:** CD

**Explanation:**

When loading data into Snowflake, it is recommended to:

? C. Load files that are approximately 100-250 MB (or larger): This size is optimal for parallel processing and can help to maximize throughput. Smaller files can lead to overhead that outweighs the actual data processing time.

? D. Avoid using embedded characters such as commas for numeric data types:

Embedded characters can cause issues during data loading as they may be interpreted incorrectly. It's best to clean the data of such characters to ensure accurate and efficient data loading.

These best practices are designed to optimize the data loading process, ensuring that data is loaded quickly and accurately into Snowflake.

References:

? Snowflake Documentation on Data Loading Considerations

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

**NEW QUESTION 88**

- (Topic 1)

Which command can be used to stage local files from which Snowflake interface?

- A. SnowSQL
- B. Snowflake classic web interface (UI)
- C. Snowsight

D. .NET driver

**Answer:** A

**Explanation:**

SnowSQL is the command-line client for Snowflake that allows users to execute SQL queries and perform all DDL and DML operations, including staging files for bulk data loading. It is specifically designed for scripting and automating tasks. References:

? SnowPro Core Certification Exam Study Guide

? Snowflake Documentation on SnowSQL <https://docs.snowflake.com/en/user-guide/snowsqli-use.html>

**NEW QUESTION 90**

- (Topic 1)

Which Snowflake partner specializes in data catalog solutions?

- A. Alation
- B. DataRobot
- C. dbt
- D. Tableau

**Answer:** A

**Explanation:**

Alation is known for specializing in data catalog solutions and is a partner of Snowflake. Data catalog solutions are essential for organizations to effectively manage their metadata and make it easily accessible and understandable for users, which aligns with the capabilities provided by Alation.

References:

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

? Snowflake??s official documentation and partner listings

**NEW QUESTION 92**

- (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 operations3.

**NEW QUESTION 94**

- (Topic 2)

Which file formats are supported for unloading data from Snowflake? (Choose two.)

- A. Avro
- B. JSON
- C. ORC
- D. XML
- E. Delimited (CSV, TSV, etc.)

**Answer:** BE

**Explanation:**

Snowflake supports unloading data in JSON and delimited file formats such as CSV and TSV. These formats are commonly used for data interchange and are supported by Snowflake for unloading operations

**NEW QUESTION 97**

- (Topic 2)

A user is preparing to load data from an external stage

Which practice will provide the MOST efficient loading performance?

- A. Organize files into logical paths
- B. Store the files on the external stage to ensure caching is maintained
- C. Use pattern matching for regular expression execution
- D. Load the data in one large file

**Answer:** A

**Explanation:**

Organizing files into logical paths can significantly improve the efficiency of data loading from an external stage. This practice helps in managing and locating files easily, which can be particularly beneficial when dealing with large datasets or complex directory structures1.

**NEW QUESTION 98**

- (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 103**

- (Topic 2)

When should a multi-cluster warehouse be used in auto-scaling mode?

- A. When it is unknown how much compute power is needed
- B. If the select statement contains a large number of temporary tables or Common Table Expressions (CTEs)
- C. If the runtime of the executed query is very slow
- D. When a large number of concurrent queries are run on the same warehouse

**Answer:** D

**Explanation:**

A multi-cluster warehouse should be used in auto-scaling mode when there is a need to handle a large number of concurrent queries. Auto-scaling allows Snowflake to automatically add or remove compute clusters to balance the load, ensuring that performance remains consistent during varying levels of demand

**NEW QUESTION 106**

- (Topic 2)

Which tasks are performed in the Snowflake Cloud Services layer? (Choose two.)

- A. Management of metadata
- B. Computing the data
- C. Maintaining Availability Zones
- D. Infrastructure security
- E. Parsing and optimizing queries

**Answer:** AE

**Explanation:**

The Snowflake Cloud Services layer performs a variety of tasks, including the management of metadata and the parsing and optimization of queries. This layer is responsible for coordinating activities across Snowflake, including user session management, security, and query compilation<sup>3</sup>.

**NEW QUESTION 110**

- (Topic 2)

Which of the following describes a Snowflake stored procedure?

- A. They can be created as secure and hide the underlying metadata from the user.
- B. They can only access tables from a single database.
- C. They can contain only a single SQL statement.
- D. They can be created to run with a caller's rights or an owner's rights.

**Answer:** D

**Explanation:**

Snowflake stored procedures can be created to execute with the privileges of the role that owns the procedure (owner's rights) or with the privileges of the role that calls the procedure (caller's rights). This allows for flexibility in managing security and access control within Snowflake<sup>1</sup>.

**NEW QUESTION 115**

- (Topic 2)

The Snowflake Cloud Data Platform is described as having which of the following architectures?

- A. Shared-disk
- B. Shared-nothing
- C. Multi-cluster shared data
- D. Serverless query engine

**Answer:** C

**Explanation:**

Snowflake's architecture is described as a multi-cluster, shared data architecture. This design combines the simplicity of a shared-disk architecture with the performance and scale-out benefits of a shared-nothing architecture, using a central repository accessible from all compute nodes<sup>2</sup>.  
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

#### NEW QUESTION 116

- (Topic 2)

What COPY INTO SQL command should be used to unload data into multiple files?

- A. SINGLE=TRUE
- B. MULTIPLE=TRUE
- C. MULTIPLE=FALSE
- D. SINGLE=FALSE

**Answer:** D

#### Explanation:

The COPY INTO SQL command with the option SINGLE=FALSE is used to unload data into multiple files. This option allows the data to be split into multiple files during the unload process. References: SnowPro Core Certification COPY INTO SQL command unload multiple files

#### NEW QUESTION 118

- (Topic 2)

When cloning a database containing stored procedures and regular views, that have fully qualified table references, which of the following will occur?

- A. The cloned views and the stored procedures will reference the cloned tables in the cloned database.
- B. An error will occur, as views with qualified references cannot be cloned.
- C. An error will occur, as stored objects cannot be cloned.
- D. The stored procedures and views will refer to tables in the source database.

**Answer:** A

#### Explanation:

When cloning a database containing stored procedures and regular views with fully qualified table references, the cloned views and stored procedures will reference the cloned tables in the cloned database (A). This ensures that the cloned database is a self-contained copy of the original, with all references pointing to objects within the same cloned database. References: SnowPro Core Certification cloning database stored procedures views

#### NEW QUESTION 123

- (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 124

- (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 126

- (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 131

- (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 135

- (Topic 2)

When loading data into Snowflake, how should the data be organized?

- A. Into single files with 100-250 MB of compressed data per file
- B. Into single files with 1-100 MB of compressed data per file
- C. Into files of maximum size of 1 GB of compressed data per file
- D. Into files of maximum size of 4 GB of compressed data per file

**Answer:** A

#### Explanation:

When loading data into Snowflake, it is recommended to organize the data into single files with 100-250 MB of compressed data per file. This size range is optimal for parallel processing and can help in achieving better performance during data loading operations. References: [COF-C02] SnowPro Core Certification Exam Study Guide

#### NEW QUESTION 138

- (Topic 2)

Which of the following is an example of an operation that can be completed without requiring compute, assuming no queries have been executed previously?

- A. SELECT SUM (ORDER\_AMT) FROM SALES;
- B. SELECT AVG(ORDER\_QTY) FROM SALES;
- C. SELECT MIN(ORDER\_AMT) FROM SALES;
- D. SELECT ORDER\_AMT \* ORDER\_QTY FROM SALES;

**Answer:** B

#### Explanation:

Operations that do not require compute resources are typically those that can leverage previously cached results. However, if no queries have been executed previously, all the given operations would require compute to execute. It's important to note that certain operations like DDL statements and queries that hit the result cache do not consume compute credits<sup>2</sup>.

#### NEW QUESTION 139

- (Topic 2)

Network policies can be set at which Snowflake levels? (Choose two.)

- A. Role
- B. Schema
- C. User
- D. Database
- E. Account
- F. Tables

**Answer:** CE

#### Explanation:

Network policies in Snowflake can be set at the user level and at the account level<sup>2</sup>.

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

#### NEW QUESTION 144

- (Topic 2)

How can a row access policy be applied to a table or a view? (Choose two.)

- A. Within the policy DDL
- B. Within the create table or create view DDL
- C. By future APPLY for all objects in a schema
- D. Within a control table
- E. Using the command ALTER <object> ADD ROW ACCESS POLICY <policy>;

**Answer:** AE

**Explanation:**

A row access policy can be applied to a table or a view within the policy DDL when defining the policy. Additionally, an existing row access policy can be applied to a table or a view using the ALTER <object> ADD ROW ACCESS POLICY <policy> command

**NEW QUESTION 146**

- (Topic 2)

What type of query benefits the MOST from search optimization?

- A. A query that uses only disjunction (i.e., OR) predicates
- B. A query that includes analytical expressions
- C. A query that uses equality predicates or predicates that use IN
- D. A query that filters on semi-structured data types

**Answer:** C

**Explanation:**

Search optimization in Snowflake is designed to improve the performance of queries that are selective and involve point lookup operations using equality and IN predicates. It is particularly beneficial for queries that access columns with a high number of distinct values<sup>1</sup>.

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

**NEW QUESTION 150**

- (Topic 2)

What is the minimum Snowflake edition required for row level security?

- A. Standard
- B. Enterprise
- C. Business Critical
- D. Virtual Private Snowflake

**Answer:** B

**Explanation:**

Row level security in Snowflake is available starting with the Enterprise edition. This feature allows for the creation of row access policies that can control access to data at the row level within tables and views

**NEW QUESTION 155**

- (Topic 2)

Which snowflake objects will incur both storage and cloud compute charges? (Select TWO)

- A. Materialized view
- B. Sequence
- C. Secure view
- D. Transient table
- E. Clustered table

**Answer:** AD

**Explanation:**

In Snowflake, both materialized views and transient tables will incur storage charges because they store data. They will also incur compute charges when queries are run against them, as compute resources are used to process the queries. References:

[COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 158**

- (Topic 2)

What is true about sharing data in Snowflake? (Choose two.)

- A. The Data Consumer pays for data storage as well as for data computing.
- B. The shared data is copied into the Data Consumer account, so the Consumer can modify it without impacting the base data of the Provider.
- C. A Snowflake account can both provide and consume shared data.
- D. The Provider is charged for compute resources used by the Data Consumer to query the shared data.
- E. The Data Consumer pays only for compute resources to query the shared data.

**Answer:** CE

**Explanation:**

In Snowflake's data sharing model, any full Snowflake account can both provide and consume shared data. Additionally, the data consumer pays only for the compute resources used to query the shared data. No actual data is copied or transferred between accounts, and shared data does not take up any storage in a consumer account, so the consumer does not pay for data storage<sup>1</sup>.

References = Introduction to Secure Data Sharing | Snowflake Documentation

#### NEW QUESTION 162

- (Topic 2)

Which of the following accurately describes shares?

- A. Tables, secure views, and secure UDFs can be shared
- B. Shares can be shared
- C. Data consumers can clone a new table from a share
- D. Access to a share cannot be revoked once granted

**Answer:** A

#### Explanation:

Shares in Snowflake are named objects that encapsulate all the information required to share databases, schemas, tables, secure views, and secure UDFs. These objects can be added to a share by granting privileges on them to the share via a database role

#### NEW QUESTION 165

- (Topic 2)

In the Snowflake access control model, which entity owns an object by default?

- A. The user who created the object
- B. The SYSADMIN role
- C. Ownership depends on the type of object
- D. The role used to create the object

**Answer:** D

#### Explanation:

In Snowflake's access control model, the default owner of an object is the role that was used to create the object. This role has the OWNERSHIP privilege on the object and can grant access to other roles<sup>1</sup>

#### NEW QUESTION 169

- (Topic 2)

A running virtual warehouse is suspended.

What is the MINIMUM amount of time that the warehouse will incur charges for when it is restarted?

- A. 1 second
- B. 60 seconds
- C. 5 minutes
- D. 60 minutes

**Answer:** B

#### Explanation:

When a running virtual warehouse in Snowflake is suspended and then restarted, the minimum amount of time it will incur charges for is 60 seconds<sup>2</sup>.

#### NEW QUESTION 171

- (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 172

- (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<sup>2</sup>.

#### NEW QUESTION 177

- (Topic 2)

How long is the Fail-safe period for temporary and transient tables?

- A. There is no Fail-safe period for these tables.
- B. 1 day
- C. 7 days
- D. 31 days
- E. 90 days

**Answer:** A

**Explanation:**

Temporary and transient tables in Snowflake do not have a Fail-safe period. Once the session ends or the tables are dropped, the data is purged and not recoverable<sup>1</sup>.

Reference: <https://docs.snowflake.com/en/user-guide/tables-temp-transient.html>

**NEW QUESTION 179**

- (Topic 3)

Which SQL command can be used to see the CREATE definition of a masking policy?

- A. SHOW MASKING POLICIES
- B. DESCRIBE MASKING POLICY
- C. GET\_DDL
- D. LIST MASKING POLICIES

**Answer:** C

**Explanation:**

The SQL command GET\_DDL can be used to retrieve the CREATE definition of a masking policy in Snowflake. This command generates the DDL statement required to recreate the masking policy

**NEW QUESTION 180**

- (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 Documentation<sup>1</sup>.

**NEW QUESTION 181**

- (Topic 3)

Which TABLE function helps to convert semi-structured data to a relational representation?

- A. CHECK\_JSON
- B. TO\_JSON
- C. FLATTEN
- D. PARSE\_JSON

**Answer:** C

**Explanation:**

The FLATTEN table function in Snowflake is used to convert semi- structured data, such as JSON or XML, into a relational format. It expands nested arrays or objects into multiple rows, making the data suitable for relational querying<sup>3</sup>.

**NEW QUESTION 184**

- (Topic 3)

How does Snowflake allow a data provider with an Azure account in central Canada to share data with a data consumer on AWS in Australia?

- A. The data provider in Azure Central Canada can create a direct share to AWS Asia Pacific, if they are both in the same organization.
- B. The data consumer and data provider can form a Data Exchange within the same organization to create a share from Azure Central Canada to AWS Asia Pacific.
- C. The data provider uses the GET DATA workflow in the Snowflake Data Marketplace to create a share between Azure Central Canada and AWS Asia Pacific.
- D. The data provider must replicate the database to a secondary account in AWS Asia Pacific within the same organization then create a share to the data consumer's account.

**Answer:** D

**Explanation:**

Snowflake allows data providers to share data with consumers across different cloud platforms and regions through database replication. The data provider must

replicate the database to a secondary account in the target region or cloud platform within the same organization, and then create a share to the data consumer's account. This process ensures that the data is available in the consumer's region and on their cloud platform, facilitating seamless data sharing. References: Sharing data securely across regions and cloud platforms | Snowflake Documentation

#### NEW QUESTION 186

- (Topic 3)

Which role has the ability to create and manage users and roles?

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

**Answer:** B

#### Explanation:

The USERADMIN role in Snowflake has the ability to create and manage users and roles within the Snowflake environment. This role is specifically dedicated to user and role management and creation

#### NEW QUESTION 187

- (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 190

- (Topic 3)

Which of the following are considerations when using a directory table when working with unstructured data? (Choose two.)

- A. A directory table is a separate database object.
- B. Directory tables store data file metadata.
- C. A directory table will be automatically added to a stage.
- D. Directory tables do not have their own grantable privileges.
- E. Directory table data can not be refreshed manually.

**Answer:** BD

#### Explanation:

Directory tables in Snowflake are used to store metadata about data files in a stage. They are not separate database objects but are conceptually similar to external tables. Directory tables do not have grantable privileges of their own

#### NEW QUESTION 191

- (Topic 3)

What effect does WAIT\_FOR\_COMPLETION = TRUE have when running an ALTER WAREHOUSE command and changing the warehouse size?

- A. The warehouse size does not change until all queries currently running in the warehouse have completed.
- B. The warehouse size does not change until all queries currently in the warehouse queue have completed.
- C. The warehouse size does not change until the warehouse is suspended and restarted.
- D. It does not return from the command until the warehouse has finished changing its size.

**Answer:** D

#### Explanation:

The WAIT\_FOR\_COMPLETION = TRUE parameter in an ALTER WAREHOUSE command ensures that the command does not return until the warehouse has completed resizing. This means that the command will wait until all the necessary compute resources have been provisioned and the warehouse size has been changed. References: [COF-C02] SnowPro Core Certification Exam Study Guide

#### NEW QUESTION 194

- (Topic 3)

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;GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO ROLE MYROLE;
- B. GRANT ROLE MYROLE TO USER USER1;GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER USER1;
- C. GRANT ROLE MYROLE TO USER USER1;GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYD
- D. K"-SCHEMA TO USER! ;
- E. GRANT ROLE MYROLE TO USER USER1;GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO MYROLE;

**Answer:** A

**Explanation:**

To provide a user with the necessary access to create a materialized view in a schema, the user must be granted a role that has the CREATE MATERIALIZED VIEW privilege on that schema. First, the role is granted to the user, and then the privilege is granted to the role

**NEW QUESTION 195**

- (Topic 3)

Which type of join will list all rows in the specified table, even if those rows have no match in the other table?

- A. Cross join
- B. Inner join
- C. Natural join
- D. Outer join

**Answer:** D

**Explanation:**

An outer join, specifically a left outer join, will list all rows from the left table and match them with rows from the right table. If there is no match, the result will still include the row from the left table, with NULLs for columns from the right table. References: Based on general SQL knowledge as of 2021.

**NEW QUESTION 199**

- (Topic 3)

Which objects together comprise a namespace in Snowflake? (Select TWO).

- A. Account
- B. Database
- C. Schema
- D. Table
- E. Virtual warehouse

**Answer:** BC

**Explanation:**

In Snowflake, a namespace is comprised of a database and a schema. The combination of a database and schema uniquely identifies database objects within an account

**NEW QUESTION 200**

- (Topic 3)

What is the recommended way to change the existing file format type in my format from CSV to JSON?

- A. ALTER FILE FORMAT my\_format SET TYPE=JSON;
- B. ALTER FILE FORMAT my format SWAP TYPE WITH JSON;
- C. CREATE OR REPLACE FILE FORMAT my format TYPE=JSON;
- D. REPLACE FILE FORMAT my format TYPE=JSON;

**Answer:** A

**Explanation:**

To change the existing file format type from CSV to JSON, the recommended way is to use the ALTER FILE FORMAT command with the SET TYPE=JSON clause. This alters the file format specification to use JSON instead of CSV. References: Based on my internal knowledge as of 2021.

**NEW QUESTION 202**

- (Topic 3)

Which formats does Snowflake store unstructured data in? (Choose two.)

- A. GeoJSON
- B. Array
- C. XML
- D. Object
- E. BLOB

**Answer:** AC

**Explanation:**

Snowflake supports storing unstructured data and provides native support for semi-structured file formats such as JSON, Avro, Parquet, ORC, and XML1. GeoJSON, being a type of JSON, and XML are among the formats that can be stored in Snowflake. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 204**

- (Topic 3)

Which features could be used to improve the performance of queries that return a small subset of rows from a large table? (Select TWO).

- A. Search optimization service
- B. Automatic clustering
- C. Row access policies
- D. Multi-cluster virtual warehouses
- E. Secure views

**Answer:** AB

**Explanation:**

The search optimization service and automatic clustering are features that can improve the performance of queries returning a small subset of rows from a large table. The search optimization service is designed for low-latency point lookup queries, while automatic clustering organizes data in micro-partitions based on specific dimensions to reduce the amount of data scanned during queries.

**NEW QUESTION 206**

- (Topic 3)

What happens when a database is cloned?

- A. It does not retain any privileges granted on the source object.
- B. It replicates all granted privileges on the corresponding source objects.
- C. It replicates all granted privileges on the corresponding child objects.
- D. It replicates all granted privileges on the corresponding child schema objects.

**Answer:** A

**Explanation:**

When a database is cloned in Snowflake, it does not retain any privileges that were granted on the source object. The clone will need to have privileges reassigned as necessary for users to access it. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 208**

- (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 year1.

**NEW QUESTION 210**

- (Topic 3)

What service is provided as an integrated Snowflake feature to enhance Multi-Factor Authentication (MFA) support?

- A. Duo Security
- B. OAuth
- C. Okta
- D. Single Sign-On (SSO)

**Answer:** A

**Explanation:**

Snowflake provides Multi-Factor Authentication (MFA) support as an integrated feature, powered by the Duo Security service. This service is managed completely by Snowflake, and users do not need to sign up separately with Duo1

**NEW QUESTION 212**

- (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 file3.

**NEW QUESTION 217**

- (Topic 3)

Which Snowflake object can be accessed in the FROM clause of a query, returning a set of rows having one or more columns?

- A. A User-Defined Table Function (UDTF)
- B. A Scalar User Function (UDF)
- C. A stored procedure
- D. A task

**Answer:** A

**Explanation:**

In Snowflake, a User-Defined Table Function (UDTF) can be accessed in the FROM clause of a query. UDTFs return a set of rows with one or more columns, which can be queried like a regular table

**NEW QUESTION 221**

- (Topic 3)

Which of the following are characteristics of security in Snowflake?

- A. Account and user authentication is only available with the Snowflake Business Critical edition.
- B. Support for HIPAA and GDPR compliance is available for UI Snowflake editions.
- C. Periodic rekeying of encrypted data is available with the Snowflake Enterprise edition and higher
- D. Private communication to internal stages is allowed in the Snowflake Enterprise edition and higher.

**Answer:** C

**Explanation:**

One of the security features of Snowflake includes the periodic rekeying of encrypted data, which is available with the Snowflake Enterprise edition and higher<sup>2</sup>. This ensures that the encryption keys are rotated regularly to maintain a high level of security. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 225**

- (Topic 3)

What is the MINIMUM Snowflake edition required to use the periodic rekeying of micro- partitions?

- A. Enterprise
- B. Business Critical
- C. Standard
- D. Virtual Private Snowflake

**Answer:** A

**Explanation:**

Periodic rekeying of micro-partitions is a feature that requires the Enterprise Edition of Snowflake or higher. This feature is part of Snowflake's comprehensive approach to encryption key management, ensuring data security through regular rekeying<sup>1</sup>. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 228**

- (Topic 3)

Snowflake's hierarchical key mode includes which keys? (Select TWO).

- A. Account master keys
- B. Database master keys
- C. File keys
- D. Secure view keys
- E. Schema master keys

**Answer:** AC

**Explanation:**

Snowflake's hierarchical key model includes several levels of keys, where Account master keys and File keys are part of this hierarchy. Account master keys are used to encrypt all the data within an account, while File keys are used to encrypt individual files within the database<sup>2</sup>.

**NEW QUESTION 233**

- (Topic 3)

What are advantages clones have over tables created with CREATE TABLE AS SELECT statement? (Choose two.)

- A. The clone always stays in sync with the original table.
- B. The clone has better query performance.
- C. The clone is created almost instantly.
- D. The clone will have time travel history from the original table.
- E. The clone saves space by not duplicating storage.

**Answer:** CE

**Explanation:**

Clones in Snowflake have the advantage of being created almost instantly and saving space by not duplicating storage. This is due to Snowflake's zero-copy cloning feature, which allows for the creation of object clones without the additional storage costs typically associated with data duplication<sup>23</sup>. Clones are independent of the original table and do not stay in sync with it, nor do they inherently have better query performance. However, they do inherit the time travel history from the original table at the time of cloning

**NEW QUESTION 235**

- (Topic 3)

What is a characteristic of the Snowflake Query Profile?

- A. It can provide statistics on a maximum number of 100 queries per week.
- B. It provides a graphic representation of the main components of the query processing.
- C. It provides detailed statistics about which queries are using the greatest number of compute resources.

D. It can be used by third-party software using the Query Profile API.

**Answer:** B

**Explanation:**

The Snowflake Query Profile provides a graphic representation of the main components of the query processing. This visual aid helps users understand the execution details and performance characteristics of their queries<sup>4</sup>.

**NEW QUESTION 240**

- (Topic 3)

What can a Snowflake user do with the information included in the details section of a Query Profile?

- A. Determine the total duration of the query.
- B. Determine the role of the user who ran the query.
- C. Determine the source system that the queried table is from.
- D. Determine if the query was on structured or semi-structured data.

**Answer:** A

**Explanation:**

The details section of a Query Profile in Snowflake provides users with various statistics and information about the execution of a query. One of the key pieces of information that can be determined from this section is the total duration of the query, which helps in understanding the performance and identifying potential bottlenecks. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 245**

- (Topic 3)

Which stages are used with the Snowflake PUT command to upload files from a local file system? (Choose three.)

- A. Schema Stage
- B. User Stage
- C. Database Stage
- D. Table Stage
- E. External Named Stage
- F. Internal Named Stage

**Answer:** BDF

**Explanation:**

The Snowflake PUT command is used to upload files from a local file system to Snowflake stages, specifically the user stage, table stage, and internal named stage. These stages are where the data files are temporarily stored before being loaded into Snowflake tables

**NEW QUESTION 250**

- (Topic 3)

What is the MAXIMUM Time Travel retention period for a transient table?

- A. 0 days
- B. 1 day
- C. 7 days
- D. 90 days

**Answer:** B

**Explanation:**

The maximum Time Travel retention period for a transient table in Snowflake is 1 day. This is the default and maximum duration for which Snowflake maintains the historical data for transient tables, allowing users to query data as it appeared at any point within the past 24 hours<sup>2</sup>.

**NEW QUESTION 252**

- (Topic 3)

A company needs to read multiple terabytes of data for an initial load as part of a Snowflake migration. The company can control the number and size of CSV extract files.

How does Snowflake recommend maximizing the load performance?

- A. Use auto-ingest Snowpipes to load large files in a serverless model.
- B. Produce the largest files possible, reducing the overall number of files to process.
- C. Produce a larger number of smaller files and process the ingestion with size Small virtual warehouses.
- D. Use an external tool to issue batched row-by-row inserts within BEGIN TRANSACTION and COMMIT commands.

**Answer:** B

**Explanation:**

Snowflake's documentation recommends producing the largest files possible for data loading, as larger files reduce the number of files to process and the overhead associated with handling many small files. This approach can maximize the load performance by leveraging Snowflake's ability to ingest large files efficiently<sup>1</sup>. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 256**

- (Topic 3)

Which feature allows a user the ability to control the organization of data in a micro- partition?

- A. Range Partitioning
- B. Search Optimization Service
- C. Automatic Clustering
- D. Horizontal Partitioning

**Answer:** C

**Explanation:**

Automatic Clustering is a feature that allows users to control the organization of data within micro-partitions in Snowflake. By defining clustering keys, Snowflake can automatically reorganize the data in micro-partitions to optimize query performance<sup>1</sup>.

**NEW QUESTION 261**

- (Topic 3)

When would Snowsight automatically detect if a target account is in a different region and enable cross-cloud auto-fulfillment?

- A. When using a paid listing on the Snowflake Marketplace
- B. When using a private listing on the Snowflake Marketplace
- C. When using a personalized listing on the Snowflake Marketplace
- D. When using a Direct Share with another account

**Answer:** A

**Explanation:**

Snowsight automatically detects if a target account is in a different region and enables cross-cloud auto-fulfillment when using a paid listing on the Snowflake Marketplace. This feature allows Snowflake to manage the replication of data products to consumer regions as needed, without manual intervention<sup>1</sup>.

**NEW QUESTION 264**

- (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 Snowflake<sup>2</sup>. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 265**

- (Topic 3)

If file format options are specified in multiple locations, the load operation selects which option FIRST to apply in order of precedence?

- A. Table definition
- B. Stage definition
- C. Session level
- D. COPY INTO TABLE statement

**Answer:** D

**Explanation:**

When file format options are specified in multiple locations, the load operation applies the options in the following order of precedence: first, the COPY INTO TABLE statement; second, the stage definition; and third, the table definition<sup>1</sup>

**NEW QUESTION 267**

- (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 270**

- (Topic 3)

For non-materialized views, what column in Information Schema and Account Usage identifies whether a view is secure or not?

- A. CHECK\_OPTION
- B. IS\_SECURE
- C. IS\_UPDATEABLE
- D. TABLE\_NAME

**Answer:** B

**Explanation:**

In the Information Schema and Account Usage, the column that identifies whether a view is secure or not is IS\_SECURE2.

**NEW QUESTION 271**

- (Topic 3)

The first user assigned to a new account, ACCOUNTADMIN, should create at least one additional user with which administrative privilege?

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

**Answer:** A

**Explanation:**

The first user assigned to a new Snowflake account, typically with the ACCOUNTADMIN role, should create at least one additional user with the USERADMIN administrative privilege. This role is responsible for creating and managing users and roles within the Snowflake account. References: Access control considerations | Snowflake Documentation

**NEW QUESTION 274**

- (Topic 3)

If queries start to queue in a multi-cluster virtual warehouse, an additional compute cluster starts immediately under what setting?

- A. Auto-scale mode
- B. Maximized mode
- C. Economy scaling policy
- D. Standard scaling policy

**Answer:** A

**Explanation:**

In Snowflake, when queries begin to queue in a multi-cluster virtual warehouse, an additional compute cluster starts immediately if the warehouse is set to auto-scale mode. This mode allows Snowflake to automatically add or resume additional clusters as soon as the workload increases, and similarly, shut down or pause the additional clusters when the load decreases

**NEW QUESTION 275**

- (Topic 3)

Which statement describes how Snowflake supports reader accounts?

- A. A reader account can consume data from the provider account that created it and combine it with its own data.
- B. A consumer needs to become a licensed Snowflake customer as data sharing is only supported between Snowflake accounts.
- C. The users in a reader account can query data that has been shared with the reader account and can perform DML tasks.
- D. The SHOW MANAGED ACCOUNTS command will view all the reader accounts that have been created for an account.

**Answer:** B

**Explanation:**

Snowflake supports reader accounts, which are a type of account that allows data providers to share data with consumers who are not Snowflake customers. However, for data sharing to occur, the consumer needs to become a licensed Snowflake customer because data sharing is only supported between Snowflake accounts. References: Introduction to Secure Data Sharing | Snowflake Documentation2.

**NEW QUESTION 278**

- (Topic 3)

What does Snowflake's search optimization service support?

- A. External tables
- B. Materialized views
- C. Tables and views that are not protected by row access policies
- D. Casts on table columns (except for fixed-point numbers cast to strings)

**Answer:** C

**Explanation:**

Snowflake's search optimization service supports tables and views that are not protected by row access policies. It is designed to improve the performance of certain types of queries on tables, including selective point lookup queries and queries on fields in VARIANT, OBJECT, and ARRAY (semi-structured) columns1.

**NEW QUESTION 281**

- (Topic 3)

When unloading data to an external stage, what is the MAXIMUM file size supported?

- A. 1 GB
- B. 5 GB
- C. 10 GB
- D. 16 GB

**Answer:** B

**Explanation:**

When unloading data to an external stage, the maximum file size supported is 5 GB. This limit ensures efficient data transfer and management within Snowflake's architecture

**NEW QUESTION 283**

- (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 285**

- (Topic 3)

Two users share a virtual warehouse named wh dev 01. When one of the users loads data, the other one experiences performance issues while querying data. How does Snowflake recommend resolving this issue?

- A. Scale up the existing warehouse.
- B. Create separate warehouses for each user.
- C. Create separate warehouses for each workload.
- D. Stop loading and querying data at the same time.

**Answer:** C

**Explanation:**

Snowflake recommends creating separate warehouses for each workload to resolve performance issues caused by shared virtual warehouses. This ensures that the resources are not being overutilized by one user's activities, thereby affecting the performance of another user's activities<sup>4</sup>.

**NEW QUESTION 287**

- (Topic 3)

What role is required to use Partner Connect?

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

**Answer:** A

**Explanation:**

To use Partner Connect, the ACCOUNTADMIN role is required. Partner Connect allows account administrators to easily create trial accounts with selected Snowflake business partners and integrate these accounts with Snowflake

**NEW QUESTION 288**

- (Topic 3)

How many resource monitors can be assigned at the account level?

- A. 1
- B. 2
- C. 3
- D. 4

**Answer:** A

**Explanation:**

Snowflake allows for only one resource monitor to be assigned at the account level. This monitor oversees the credit usage of all the warehouses in the account. References: Snowflake Documentation

**NEW QUESTION 291**

- (Topic 3)

A user has a standard multi-cluster warehouse auto-scaling policy in place.

Which condition will trigger a cluster to shut-down?

- A. When after 2-3 consecutive checks the system determines that the load on the most- loaded cluster could be redistributed.
- B. When after 5-6 consecutive checks the system determines that the load on the most- loaded cluster could be redistributed.
- C. When after 5-6 consecutive checks the system determines that the load on the least- loaded cluster could be redistributed.
- D. When after 2-3 consecutive checks the system determines that the load on the least- loaded cluster could be redistributed.

**Answer:** D

**Explanation:**

In a standard multi-cluster warehouse with auto-scaling, a cluster will shut down when, after 2-3 consecutive checks, the system determines that the load on the least-loaded cluster could be redistributed to other clusters. This ensures efficient resource utilization and cost management. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 292**

- (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 c,their%20network%20connection%20to%20Snowflake.](https://docs.snowflake.com/en/user-guide/snowcd.html#:~:text=SnowCD%20(i.e.%20Snowflake%20Connectivity%20Diagnosti c,their%20network%20connection%20to%20Snowflake.)

**NEW QUESTION 296**

- (Topic 3)

Which statement accurately describes a characteristic of a materialized view?

- A. A materialized view can query only a single table.
- B. Data accessed through materialized views can be stale.
- C. Materialized view refreshes need to be maintained by the user.
- D. Querying a materialized view is slower than executing a query against the base table of the view.

**Answer:** B

**Explanation:**

A characteristic of a materialized view is that the data accessed through it can be stale. This is because the data in a materialized view may not reflect the latest changes in the base tables until the view is refreshed

**NEW QUESTION 298**

- (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 UDF2.

**NEW QUESTION 300**

- (Topic 3)

Which REST API can be used with unstructured data?

- A. inscrtFilcs
- B. insertReport
- C. GET /api/tiles/
- D. loadHistoryScan

**Answer:** C

**Explanation:**

The REST API used with unstructured data in Snowflake is GET /api/files/, which retrieves (downloads) a data file from an internal or external stage4.

**NEW QUESTION 303**

- (Topic 4)

Which commands can only be executed using SnowSQL? (Select TWO).

- A. COPY INTO
- B. GET
- C. LIST
- D. PUT
- E. REMOVE

**Answer:** CD

**Explanation:**

The LIST and PUT commands are specific to SnowSQL and cannot be executed in the web interface or other SQL clients. LIST is used to display the contents of a stage, and PUT is used to upload files to a stage. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 306**

- (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 309**

- (Topic 4)

When enabling access to unstructured data, which URL permits temporary access to a staged file without the need to grant privileges to the stage or to issue access tokens?

- A. File URL
- B. Scoped URL
- C. Relative URL
- D. Pre-Signed URL

**Answer:** B

**Explanation:**

A Scoped URL permits temporary access to a staged file without the need to grant privileges to the stage or to issue access tokens. It provides a secure way to share access to files stored in Snowflake

**NEW QUESTION 311**

- (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 314**

- (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 319

- (Topic 4)

What metadata does Snowflake store for rows in micro-partitions? (Select TWO).

- A. Range of values
- B. Distinct values
- C. Index values
- D. Sorted values
- E. Null values

**Answer:** AB

**Explanation:**

Snowflake stores metadata for rows in micro-partitions, including the range of values for each column and the number of distinct values<sup>1</sup>.

#### NEW QUESTION 322

- (Topic 4)

When working with a managed access schema, who has the OWNERSHIP privilege of any tables added to the schema?

- A. The database owner
- B. The object owner
- C. The schema owner
- D. The Snowflake user's role

**Answer:** C

**Explanation:**

In a managed access schema, the schema owner retains the OWNERSHIP privilege of any tables added to the schema. This means that while object owners have certain privileges over the objects they create, only the schema owner can manage privilege grants on these objects<sup>1</sup>.

#### NEW QUESTION 324

- (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 326

- (Topic 4)

What is the minimum Snowflake Edition that supports secure storage of Protected Health Information (PHI) data?

- A. Standard Edition
- B. Enterprise Edition
- C. Business Critical Edition
- D. Virtual Private Snowflake Edition

**Answer:** C

**Explanation:**

The minimum Snowflake Edition that supports secure storage of Protected Health Information (PHI) data is the Business Critical Edition. This edition offers enhanced security features necessary for compliance with regulations such as HIPAA and HITRUST CSF<sup>4</sup>.

#### NEW QUESTION 327

- (Topic 4)

Which object can be used with Secure Data Sharing?

- A. View
- B. Materialized view
- C. External table
- D. User-Defined Function (UDF)

**Answer:** A

**Explanation:**

Views can be used with Secure Data Sharing in Snowflake. Materialized views, external tables, and UDFs are not typically shared directly for security and performance reasons<sup>2</sup>.

#### NEW QUESTION 331

- (Topic 4)

What does a masking policy consist of in Snowflake?

- A. A single data type, with one or more conditions, and one or more masking functions
- B. A single data type, with only one condition, and only one masking function
- C. Multiple data types, with only one condition, and one or more masking functions
- D. Multiple data types, with one or more conditions, and one or more masking functions

**Answer:** A

**Explanation:**

A masking policy in Snowflake consists of a single data type, with one or more conditions, and one or more masking functions. These components define how the data is masked based on the specified conditions<sup>3</sup>.

**NEW QUESTION 332**

- (Topic 4)

What objects in Snowflake are supported by Dynamic Data Masking? (Select TWO).'

- A. Views
- B. Materialized views
- C. Tables
- D. External tables
- E. Future grants

**Answer:** AC

**Explanation:**

Dynamic Data Masking in Snowflake supports tables and views. These objects can have masking policies applied to their columns to dynamically mask data at query time<sup>3</sup>.

**NEW QUESTION 335**

- (Topic 4)

Which Snowflake table objects can be shared with other accounts? (Select TWO).

- A. Temporary tables
- B. Permanent tables
- C. Transient tables
- D. External tables
- E. User-Defined Table Functions (UDTFs)

**Answer:** BD

**Explanation:**

In Snowflake, permanent tables and external tables can be shared with other accounts using Secure Data Sharing. Temporary tables, transient tables, and UDTFs are not shareable objects

**NEW QUESTION 338**

- (Topic 4)

What does the LATERAL modifier for the FLATTEN function do?

- A. Casts the values of the flattened data
- B. Extracts the path of the flattened data
- C. Joins information outside the object with the flattened data
- D. Retrieves a single instance of a repeating element in the flattened data

**Answer:** C

**Explanation:**

The LATERAL modifier for the FLATTEN function allows joining information outside the object (such as other columns in the source table) with the flattened data, creating a lateral view that correlates with the preceding tables in the FROM clause<sup>2345</sup>. References: [COF-C02] SnowPro Core Certification Exam Study Guide

**NEW QUESTION 341**

- (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>78</sup>.

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

#### NEW QUESTION 346

- (Topic 4)

What step can reduce data spilling in Snowflake?

- A. Using a larger virtual warehouse
- B. Increasing the virtual warehouse maximum timeout limit
- C. Increasing the amount of remote storage for the virtual warehouse
- D. Using a common table expression (CTE) instead of a temporary table

**Answer:** A

#### Explanation:

To reduce data spilling in Snowflake, using a larger virtual warehouse is effective because it provides more memory and local disk space, which can accommodate larger data operations and minimize the need to spill data to disk or remote storage<sup>1</sup>. References: [COF-C02] SnowPro Core Certification Exam Study Guide

#### NEW QUESTION 348

- (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 353

- (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 356

- (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 environment<sup>4</sup>.

#### NEW QUESTION 357

- (Topic 4)

What does a Notify & Suspend action for a resource monitor do?

- A. Send an alert notification to all account users who have notifications enabled.
- B. Send an alert notification to all virtual warehouse users when thresholds over 100% have been met.
- C. Send a notification to all account administrators who have notifications enabled, and suspend all assigned warehouses after all statements being executed by the warehouses have completed.
- D. Send a notification to all account administrators who have notifications enabled, and suspend all assigned warehouses immediately, canceling any statements being executed by the warehouses.

**Answer:** C

#### Explanation:

The Notify & Suspend action for a resource monitor in Snowflake sends a notification to all account administrators who have notifications enabled and suspends all assigned warehouses. However, the suspension only occurs after all currently running statements in the warehouses have been completed<sup>1</sup>. References: [COF-

## C02] SnowPro Core Certification Exam Study Guide

### NEW QUESTION 359

- (Topic 4)

What is a characteristic of materialized views in Snowflake?

- A. Materialized views do not allow joins.
- B. Clones of materialized views can be created directly by the user.
- C. Multiple tables can be joined in the underlying query of a materialized view.
- D. Aggregate functions can be used as window functions in materialized views.

**Answer: C**

#### Explanation:

One of the characteristics of materialized views in Snowflake is that they allow multiple tables to be joined in the underlying query. This enables the pre-computation of complex queries involving joins, which can significantly improve the performance of subsequent queries that access the materialized view<sup>4</sup>.  
References: [COF-C02] SnowPro Core Certification Exam Study Guide

### NEW QUESTION 364

- (Topic 4)

Which metadata table will store the storage utilization information even for dropped tables?

- A. DATABASE\_STORAGE\_USAGE\_HISTORY
- B. TABLE\_STORAGE\_METRICS
- C. STORAGE\_DAILY\_HISTORY
- D. STAGE STORAGE USAGE HISTORY

**Answer: B**

#### Explanation:

The TABLE\_STORAGE\_METRICS metadata table stores the storage utilization information, including for tables that have been dropped but are still incurring storage costs<sup>2</sup>.

### NEW QUESTION 365

- (Topic 5)

What is the Fail-safe period for a transient table in the Snowflake Enterprise edition and higher?

- A. 0 days
- B. 1 day
- C. 7 days
- D. 14 days

**Answer: A**

#### Explanation:

The Fail-safe period for a transient table in Snowflake, regardless of the edition (including Enterprise edition and higher), is 0 days. Fail-safe is a data protection feature that provides additional retention beyond the Time Travel period for recovering data in case of accidental deletion or corruption. However, transient tables are designed for temporary or short-term use and do not benefit from the Fail-safe feature, meaning that once their Time Travel period expires, data cannot be recovered.

References:

? Snowflake Documentation: Understanding Fail-safe

### NEW QUESTION 367

- (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 369

- (Topic 5)

A user wants to add additional privileges to the system-defined roles for their virtual warehouse. How does Snowflake recommend they accomplish this?

- A. Grant the additional privileges to a custom role.
- B. Grant the additional privileges to the ACCOUNTADMIN role.
- C. Grant the additional privileges to the SYSADMIN role.
- D. Grant the additional privileges to the ORGADMIN role.

**Answer:** A

**Explanation:**

Snowflake recommends enhancing the granularity and management of privileges by creating and utilizing custom roles. When additional privileges are needed beyond those provided by the system-defined roles for a virtual warehouse or any other resource, these privileges should be granted to a custom role. This approach allows for more precise control over access rights and the ability to tailor permissions to the specific needs of different user groups or applications within the organization, while also maintaining the integrity and security model of system-defined roles.

References:

? Snowflake Documentation: Roles and Privileges

**NEW QUESTION 371**

- (Topic 5)

A Snowflake user wants to temporarily bypass a network policy by configuring the user object property MINS\_TO\_BYPASS\_NETWORK\_POLICY. What should they do?

- A. Use the SECURITYADMIN role.
- B. Use the SYSADMIN role.
- C. Use the USERADMIN role.
- D. Contact Snowflake Support.

**Answer:** C

**Explanation:**

To temporarily bypass a network policy by configuring the user object property MINS\_TO\_BYPASS\_NETWORK\_POLICY, the USERADMIN role should be used. This role has the necessary privileges to modify user properties, including setting a temporary bypass for network policies, which can be crucial for enabling access under specific circumstances without permanently altering the network security configuration. References:

? Snowflake Documentation: User Management

**NEW QUESTION 375**

- (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 377**

- (Topic 5)

What are the benefits of the replication feature in Snowflake? (Select TWO).

- A. Disaster recovery
- B. Time Travel
- C. Fail-safe
- D. Database failover and fallback
- E. Data security

**Answer:** AD

**Explanation:**

The replication feature in Snowflake provides several benefits, with disaster recovery and database failover and fallback being two of the primary advantages.

Replication allows for the continuous copying of data from one Snowflake account to another, ensuring that a secondary copy of the data is available in case of outages or disasters. This capability supports disaster recovery strategies by allowing operations to quickly switch to the replicated data in a different account or region. Additionally, it facilitates database failover and fallback procedures, ensuring business continuity and minimizing downtime.

References:

? Snowflake Documentation: Data Replication

**NEW QUESTION 382**

- (Topic 5)

What is used to denote a pre-computed data set derived from a SELECT query specification and stored for later use?

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

**Answer:** C

**Explanation:**

A materialized view in Snowflake denotes a pre-computed data set derived from a SELECT query specification and stored for later use. Unlike standard views, which dynamically compute the data each time the view is accessed, materialized views store the result of the query at the time it is executed, thereby speeding up access to the data, especially for expensive aggregations on large datasets.

References:

? Snowflake Documentation: Materialized Views

**NEW QUESTION 383**

- (Topic 5)

Which common query problems are identified by the Query Profile? (Select TWO.)

- A. Syntax error
- B. Inefficient pruning
- C. Ambiguous column names
- D. Queries too large to fit in memory
- E. Object does not exist or not authorized

**Answer:** BD

**Explanation:**

The Query Profile in Snowflake can identify common query problems, including:

? B. Inefficient pruning: This refers to the inability of a query to effectively limit the amount of data being scanned, potentially leading to suboptimal performance.

? D. Queries too large to fit in memory: This indicates that a query requires more memory than is available in the virtual warehouse, which can lead to spilling to disk and degraded performance.

The Query Profile helps diagnose these issues by providing detailed execution statistics and visualizations, aiding in query optimization and troubleshooting.

References:

? Snowflake Documentation: Query Profile Top of Form

**NEW QUESTION 386**

- (Topic 5)

How does a Snowflake stored procedure compare to a User-Defined Function (UDF)?

- A. A single executable statement can call only two stored procedure
- B. In contrast, a single SQL statement can call multiple UDFs.
- C. A single executable statement can call only one stored procedur
- D. In contrast, a single SQL statement can call multiple UDFs.
- E. A single executable statement can call multiple stored procedure
- F. In contrast, multiple SQL statements can call the same UDFs.
- G. Multiple executable statements can call more than one stored procedur
- H. In contrast, a single SQL statement can call multiple UDFs.

**Answer:** B

**Explanation:**

In Snowflake, stored procedures and User-Defined Functions (UDFs) have different invocation patterns within SQL:

? Option Bis correct: A single executable statement can call only one stored procedure due to the procedural and potentially transactional nature of stored procedures. In contrast, a single SQL statement can call multiple UDFs because UDFs are designed to operate more like functions in traditional programming, where they return a value and can be embedded within SQL queries. References: Snowflake documentation comparing the operational differences between stored procedures and UDFs.

**NEW QUESTION 390**

- (Topic 5)

What are valid sub-clauses to the OVER clause for a window function? (Select TWO).

- A. GROUP BY
- B. LIMIT
- C. ORDER BY
- D. PARTITION BY
- E. UNION ALL

**Answer:** CD

**Explanation:**

Valid sub-clauses to the OVER clause for a window function in SQL are:

? C. ORDER BY: This clause specifies the order in which the rows in a partition are processed by the window function. It is essential for functions that depend on the row order, such as ranking functions.

? D. PARTITION BY: This clause divides the result set into partitions to which the window function is applied. Each partition is processed independently of other partitions, making it crucial for functions that compute values across sets of rows that share common characteristics.

These clauses are fundamental to defining the scope and order of data over which the window function operates, enabling complex analytical computations within SQL queries. References:

? Snowflake Documentation: Window Functions

**NEW QUESTION 393**

- (Topic 5)

A user has semi-structured data to load into Snowflake but is not sure what types of operations will need to be performed on the data. Based on this situation, what type of column does Snowflake recommend be used?

- A. ARRAY
- B. OBJECT

- C. TEXT
- D. VARIANT

**Answer:** D

**Explanation:**

When dealing with semi-structured data in Snowflake, and the specific types of operations to be performed on the data are not yet determined, Snowflake recommends using the VARIANT data type. The VARIANT type is highly flexible and capable of storing data in multiple formats, including JSON, AVRO, BSON, and more, within a single column. This flexibility allows users to perform various operations on the data, including querying and manipulation of nested data structures without predefined schemas.

References:

? Snowflake Documentation: Semi-structured Data Types

**NEW QUESTION 394**

- (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 395**

- (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 399**

- (Topic 5)

Which privilege is required to use the search optimization service in Snowflake?

- A. GRANT SEARCH OPTIMIZATION ON SCHEMA <schema\_name> TO ROLE <role>
- B. GRANT SEARCH OPTIMIZATION ON DATABASE <database\_name> TO ROLE <role>
- C. GRANT ADD SEARCH OPTIMIZATION ON SCHEMA <schema\_name> TO ROLE <role>
- D. GRANT ADD SEARCH OPTIMIZATION ON DATABASE <database name> TO ROLE <role>

**Answer:** C

**Explanation:**

To utilize the search optimization service in Snowflake, the correct syntax for granting privileges to a role involves specific commands that include adding search optimization capabilities:

? Option C: GRANT ADD SEARCH OPTIMIZATION ON SCHEMA <schema\_name>

TO ROLE <role>. This command grants the specified role the ability to implement search optimization at the schema level, which is essential for enhancing search capabilities within that schema.

Options A and B do not include the correct verb "ADD," which is necessary for this specific type of grant command in Snowflake. Option D incorrectly mentions the database level, as search optimization privileges are typically configured at the schema level, not the database level. References: Snowflake documentation on the use of GRANT statements for configuring search optimization.

**NEW QUESTION 403**

- (Topic 5)

Which role has the ability to create a share from a shared database by default?

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

D. ORGADMIN

**Answer:** A

**Explanation:**

By default, the ACCOUNTADMIN role in Snowflake has the ability to create a share from a shared database. This role has the highest level of access within a Snowflake account, including the management of all aspects of the account, such as users, roles, warehouses, and databases, as well as the creation and management of shares for secure data sharing with other Snowflake accounts.

References:

? Snowflake Documentation: Roles

**NEW QUESTION 406**

- (Topic 5)

What are characteristics of reader accounts in Snowflake? (Select TWO).

- A. Reader account users cannot add new data to the account.
- B. Reader account users can share data to other reader accounts.
- C. A single reader account can consume data from multiple provider accounts.
- D. Data consumers are responsible for reader account setup and data usage costs.
- E. Reader accounts enable data consumers to access and query data shared by the provider.

**Answer:** AE

**Explanation:**

Characteristics of reader accounts in Snowflake include:

? A. Reader account users cannot add new data to the account: Reader accounts are intended for data consumption only. Users of these accounts can query and analyze the data shared with them but cannot upload or add new data to the account.

? E. Reader accounts enable data consumers to access and query data shared by the provider: One of the primary purposes of reader accounts is to allow data consumers to access and perform queries on the data shared by another Snowflake account, facilitating secure and controlled data sharing.

References:

? Snowflake Documentation: Reader Accounts

**NEW QUESTION 409**

- (Topic 5)

How are network policies defined in Snowflake?

- A. They are a set of rules that define the network routes within Snowflake.
- B. They are a set of rules that dictate how Snowflake accounts can be used between multiple users.
- C. They are a set of rules that define how data can be transferred between different Snowflake accounts within an organization.
- D. They are a set of rules that control access to Snowflake accounts by specifying the IP addresses or ranges of IP addresses that are allowed to connect to Snowflake.

**Answer:** D

**Explanation:**

Network policies in Snowflake are defined as a set of rules that manage the network-level access to Snowflake accounts. These rules specify which IP addresses or IP ranges are permitted to connect to Snowflake, enhancing the security of Snowflake accounts by preventing unauthorized access. Network policies are an essential aspect of Snowflake's security model, allowing administrators to enforce access controls based on network locations.

References:

? Snowflake Documentation: Network Policies

**NEW QUESTION 410**

- (Topic 5)

What are characteristics of transient tables in Snowflake? (Select TWO).

- A. Transient tables have a Fail-safe period of 7 days.
- B. Transient tables can be cloned to permanent tables.
- C. Transient tables persist until they are explicitly dropped.
- D. Transient tables can be altered to make them permanent tables.
- E. Transient tables have Time Travel retention periods of 0 or 1 day.

**Answer:** BC

**Explanation:**

Transient tables in Snowflake are designed for temporary or intermediate workloads with the following characteristics:

? B. Transient tables can be cloned to permanent tables: This feature allows users to create copies of transient tables for permanent use, providing flexibility in managing data lifecycles.

? C. Transient tables persist until they are explicitly dropped: Unlike temporary tables that exist for the duration of a session, transient tables remain in the database until explicitly removed by a user, offering more durability for short-term data storage needs.

References:

? Snowflake Documentation: Transient Tables

**NEW QUESTION 412**

- (Topic 5)

Which command can be used to list all the file formats for which a user has access privileges?

- A. LIST
- B. ALTER FILE FORMAT
- C. DESCRIBE FILE FORMAT

## D. SHOW FILE FORMATS

**Answer:** D

**Explanation:**

The command to list all the file formats for which a user has access privileges in Snowflake is SHOW FILE FORMATS. This command provides a list of all file formats defined in the user's current session or specified database/schema, along with details such as the name, type, and creation time of each file format. It is a valuable tool for users to understand and manage the file formats available for data loading and unloading operations.

References:

? Snowflake Documentation: SHOW FILE FORMATS

## NEW QUESTION 417

- (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 418

- (Topic 5)

What is the Fail-safe retention period for transient and temporary tables?

- A. 0 days
- B. 1 day
- C. 7 days
- D. 90 days

**Answer:** A

**Explanation:**

The Fail-safe retention period for transient and temporary tables in Snowflake is 0 days. Fail-safe is a feature designed to protect data against accidental loss or deletion by retaining historical data for a period after its Time Travel retention period expires. However, transient and temporary tables, which are designed for temporary or short-term storage and operations, do not have a Fail-safe period. Once the data is deleted or the table is dropped, it cannot be recovered.

References:

? Snowflake Documentation: Understanding Fail-safe

## NEW QUESTION 419

- (Topic 5)

When using the ALLOW\_CLIENT\_MFA\_CACHING parameter, how long is a cached Multi- Factor Authentication (MFA) token valid for?

- A. 1 hour
- B. 2 hours
- C. 4 hours
- D. 8 hours

**Answer:** C

**Explanation:**

A cached MFA token is valid for up to four hours. <https://docs.snowflake.com/en/user-guide/security-mfa#using-mfa-token-caching-to-minimize-the-number-of-prompts-during-authentication-optional>

## NEW QUESTION 424

- (Topic 5)

Which function should be used to insert JSON format string data into a VARIANT field?

- A. FLATTEN
- B. CHECK\_JSON
- C. PARSE\_JSON
- D. TO\_VARIANT

**Answer:** C

**Explanation:**

To insert JSON formatted string data into a VARIANT field in Snowflake, the correct function to use is PARSE\_JSON. The PARSE\_JSON function is specifically designed to interpret a JSON formatted string and convert it into a VARIANT type, which is Snowflake's flexible format for handling semi-structured data like JSON, XML, and Avro. This function is essential for loading and querying JSON data within Snowflake, allowing users to store and manage JSON data efficiently while preserving its structure for querying purposes. This function's usage and capabilities are detailed in the Snowflake documentation, providing users with guidance

on how to handle semi-structured data effectively within their Snowflake environments.

References:

? Snowflake Documentation: PARSE\_JSON

#### NEW QUESTION 429

- (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 430

- (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 432

- (Topic 5)

A Snowflake user is writing a User-Defined Function (UDF) that includes some unqualified object names.

How will those object names be resolved during execution?

- A. Snowflake will resolve them according to the SEARCH\_PATH parameter.
- B. Snowflake will only check the schema the UDF belongs to.
- C. Snowflake will first check the current schema, and then the schema the previous query used.
- D. Snowflake will first check the current schema, and then the PUBLIC schema of the current database.

**Answer:** D

#### Explanation:

? Object Name Resolution: When unqualified object names (e.g., table name without schema) are used in a UDF, Snowflake follows a specific hierarchy to resolve them. Here's the order:

? Note: The SEARCH\_PATH parameter influences object resolution for queries, not within UDFs.

References:

? Snowflake Documentation (Object Naming Resolution): <https://docs.snowflake.com/en/sql-reference/name-resolution.html>

#### NEW QUESTION 435

.....

## Relate Links

**100% Pass Your COF-C02 Exam with ExamBible Prep Materials**

<https://www.exambible.com/COF-C02-exam/>

## Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>