

## Snowflake.COF-C03.v2026-06-01.q489

|                                                                                                                                                                                         |                                                 |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------|
| Exam Code:                                                                                                                                                                              | COF-C03                                         |
| Exam Name:                                                                                                                                                                              | SnowPro® Core Certification 2026 Exam (COF-C03) |
| Certification Provider:                                                                                                                                                                 | Snowflake                                       |
| Free Question Number:                                                                                                                                                                   | 489                                             |
| Version:                                                                                                                                                                                | v2026-06-01                                     |
| # of views:                                                                                                                                                                             | 111                                             |
| # of Questions views:                                                                                                                                                                   | 4890                                            |
| <a href="https://www.freeqas.com/qa/Snowflake/COF-C03/Snowflake.COF-C03.v2026-06-01.q489.html">https://www.freeqas.com/qa/Snowflake/COF-C03/Snowflake.COF-C03.v2026-06-01.q489.html</a> |                                                 |

### NEW QUESTION: 1

What causes objects in a data share to become unavailable to a consumer account?

- A. The DATA\_RETENTION\_TIME\_IN\_DAYS parameter in the consumer account is set to 0.
- B. The consumer account runs the GRANT IMPORTED PRIVILEGES command on the data share every 24 hours.
- C. The objects in the data share are being deleted and the grant pattern is not re-applied systematically.
- D. The consumer account acquires the data share through a private data exchange.

**Answer: C (LEAVE A REPLY)**

Objects in a data share become unavailable to a consumer account if the objects in the data share are deleted or if the permissions on these objects are altered without re-applying the grant permissions systematically.

This is because the sharing mechanism in Snowflake relies on explicit grants of permissions on specific objects (like tables, views, or secure views) to the share. If these objects are deleted or if their permissions change without updating the share accordingly, consumers can lose access.

The DATA\_RETENTION\_TIME\_IN\_DAYS parameter does not directly affect the availability of shared objects, as it controls how long Snowflake retains historical data for time travel and does not impact data sharing permissions.

Running the GRANT IMPORTED PRIVILEGES command in the consumer account is not related to the availability of shared objects; this command is used to grant privileges on imported objects within the consumer 's account and is not a routine maintenance command that would need to be run regularly.

Acquiring a data share through a private data exchange does not inherently make objects unavailable; issues would only arise if there were problems with the share configuration or if the shared objects were deleted or had their permissions altered without re-granting access to the share.

Reference: Snowflake Documentation on Managing Access to Shared Data (<https://docs.snowflake.com/en/user-guide/data-sharing-consuming.html#managing-access-to-shared-data>)

### NEW QUESTION: 2

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** ([LEAVE A REPLY](#))

To change permissions for new objects in a managed access schema, a user should be granted the MANAGE GRANTS global privilege. This privilege allows the user to manage access control through grants on all securable objects within Snowflake2.

References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

#### NEW QUESTION: 3

What privilege is required to view the load statistics on a virtual warehouse?

- A. OPERATE
- B. USAGE
- C. MONITOR
- D. MODIFY

**Answer: B** ([LEAVE A REPLY](#))

#### NEW QUESTION: 4

A reporting dashboard uses a dedicated warehouse that automatically suspends after five minutes of inactivity. Similar queries run sporadically, with gaps exceeding five minutes. Users notice that the first query after suspension takes significantly longer than subsequent queries while the warehouse remains active.

How should query performance be improved?

- A. Set MIN\_CLUSTER\_COUNT equal to MAX\_CLUSTER\_COUNT.
- B. Increase the warehouse size to X-Large.
- C. Use ALTER WAREHOUSE to increase STATEMENT\_TIMEOUT\_IN\_SECONDS.
- D. Increase the AUTO\_SUSPEND setting.

**Answer: D** ([LEAVE A REPLY](#))

A running virtual warehouse maintains a local cache of table data. When the warehouse suspends, that cache is dropped. Because the dashboard's query gaps exceed the five-minute auto-suspend period, the warehouse repeatedly suspends, and the first query after each resume runs with a cold cache. This makes it slower than subsequent similar queries. Increasing AUTO\_SUSPEND allows the warehouse to remain active longer between dashboard requests and preserves the cache for reuse. Snowflake recommends an auto-suspend value of at least ten minutes for BI and SELECT-oriented query warehouses when cache reuse is important. This setting involves a cost trade-off because an idle but running warehouse continues consuming credits. Therefore, the selected value should reflect the actual interval between dashboard queries.

#### NEW QUESTION: 5

A global company needs to share financial reports with managers who are in different regions. Managers should only see the data from their respective regions. How can these requirements be met with the least amount of operational overhead?

- A. Create unique views for each region.
- B. Build individual reports for each manager.
- C. Use row-level security on the source data table used to generate the reports.
- D. Apply Dynamic Data Masking based on a custom role assigned to each manager.

**Answer: C ([LEAVE A REPLY](#))**

The correct answer is C. Use row-level security on the source data table used to generate the reports .

The requirement is to let managers see only rows for their respective regions. This is a row-filtering requirement, so Snowflake row-level security is the best fit.

Why C is correct:

Row-level security in Snowflake is implemented using row access policies. A row access policy can dynamically filter table rows based on the current user, current role, or mapping table logic.

Example concept:

```
CREATE ROW ACCESS POLICY region_policy
AS (region STRING) RETURNS BOOLEAN - >
region IN (
SELECT allowed_region
FROM manager_region_map
WHERE manager_role = CURRENT_ROLE()
);
```

This allows one shared table or report to serve multiple managers while enforcing region-specific visibility.

Why the other options are incorrect:

- A). Creating unique views for each region increases operational overhead as regions or access rules change.
- B). Building individual reports for each manager creates the highest operational overhead and is difficult to maintain.
- D). Dynamic Data Masking hides or transforms column values. It does not filter rows by region.

Official Snowflake documentation reference:

Snowflake documentation describes row access policies as schema-level objects that determine whether a row is visible in a query result. They are used to implement row-level security.

Reference: Snowflake Documentation - Row access policies; Snowflake Documentation - Dynamic Data Masking; SnowPro Core Study Guide - Data Protection and Governance.

**NEW QUESTION: 6**

How does Snowflake support the use of structured and semi-structured data? (Select TWO)

- A.** Structured data is supported through the use of traditional relational database tables, organized into rows and columns.
- B.** Structured and semi-structured data are stored in separate physical storage locations; users must choose which type of data to use when creating a table.
- C.** Semi-structured data is supported through the use of automatically indexed database tables.
- D.** All semi-structured data is automatically converted into structured data when it is loaded into Snowflake.
- E.** Users can load and query semi-structured data using SQL using the SQL semi-structured data handling functions.

**Answer: ([SHOW ANSWER](#))**

Snowflake supports structured data in standard tables and semi-structured data in VARIANT, OBJECT, and ARRAY columns. Users can load and query both using standard SQL functions , without schema constraints for semi- structured formats.

Reference:

Snowflake Docs: Semi-Structured Data Overview

**NEW QUESTION: 7**

When using the COPY INTO < location > command, which parameter is used to unload table rows into separate files?

- A. DETAILED\_OUTPUT
- B. PARTITION BY
- C. FILE\_FORMAT
- D. HEADER

**Answer:** ([SHOW ANSWER](#))

The correct answer is B. PARTITION BY .

The PARTITION BY option in COPY INTO < location > partitions unloaded data into separate files based on the expression specified. This is commonly used to organize unloaded files by values such as date, region, or another column.

Why B is correct:

PARTITION BY controls how unloaded rows are grouped into different output paths/files.

Example:

```
COPY INTO @my_stage/sales/  
FROM sales  
PARTITION BY (TO_DATE(sale_date))  
FILE_FORMAT = (TYPE = CSV);
```

This unloads rows into separate output paths based on the sale\_date expression.

Why the other options are incorrect:

- A). DETAILED\_OUTPUT controls the amount of unload result information returned by the command.
- C). FILE\_FORMAT controls the format of the unloaded files, such as CSV, JSON, or Parquet.
- D). HEADER controls whether column headings are included in unloaded output files.

Official Snowflake documentation reference:

Snowflake documentation for COPY INTO < location > describes PARTITION BY as the option used to partition unloaded rows into separate files based on an expression.

Reference: Snowflake Documentation - COPY INTO < location > ; Snowflake Documentation - Unloading data with partitioning; SnowPro Core Study Guide - Data Loading and Unloading.

==

### NEW QUESTION: 8

Which command line parameter value can be pre-specified as an environment variable in SnowSQL?

- A. MFA-PASSCODE
- B. OPTION
- C. VARIABLE
- D. HOST

**Answer:** ([SHOW ANSWER](#))

### NEW QUESTION: 9

How does Snowflake utilize clustering information to improve query performance?

- A. It prunes unnecessary micro-partitions based on clustering metadata.
- B. It compresses the data within micro-partitions for faster querying.
- C. It automatically allocates additional resources to improve query execution.

**D.** It organizes clustering information to speed-up data retrieval from storage

**Answer:** ([SHOW ANSWER](#))

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

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

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

References:

Snowflake Documentation on Clustering

Snowflake Documentation on Micro-partition Pruning

### NEW QUESTION: 10

Snowflake's access control framework combines which models for securing data? (Select TWO).

**A.** Attribute-based Access Control (ABAC) 1

**B.** Discretionary Access Control (DAC)

**C.** Access Control List (ACL)

**D.** Role-based Access Control (RBAC)

**E.** Rule-based Access Control (RuBAC)

**Answer:** B,D ([LEAVE A REPLY](#))

Snowflake's access control framework utilizes a combination of Discretionary Access Control (DAC) and Role-based Access Control (RBAC). DAC in Snowflake allows the object owner to grant access privileges to other roles. RBAC involves assigning roles to users and then granting privileges to those roles. Through roles, Snowflake manages which users have access to specific objects and what actions they can perform, which is central to security and governance in the Snowflake environment. References: Snowflake Documentation on Access Control,

### NEW QUESTION: 11

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

**A.** They can be set for any Snowflake Edition.

**B.** They can be applied to roles.

**C.** They restrict or enable access to specific IP addresses.

**D.** They are activated using ALTER DATABASE SQL commands.

**E.** They can only be managed using the ORGADMIN role.

**Answer:** A,C ([LEAVE A REPLY](#))

Snowflake network policies are a security feature that allows administrators to control access to Snowflake by specifying allowed and blocked IP address ranges. These policies apply to all editions of Snowflake, making them widely applicable across different Snowflake environments. They are specifically designed to restrict or enable access based on the originating IP addresses of client requests, adding an extra layer of security.

Network policies are not applied to roles but are set at the account or user level. They are not activated using ALTER DATABASE SQL commands but are managed through ALTER ACCOUNT or ALTER NETWORK POLICY commands. The management of network policies does not exclusively require the ORGADMIN role; instead, they can be managed by users with the necessary privileges on the account.

**NEW QUESTION: 12**

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: C,D** ([LEAVE A REPLY](#))

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: 13**

Which view in SNOWFLAKE.ACCOUNT\_USAGE shows from which IP address a user connected to Snowflake?

- A. ACCESS\_HISTORY
- B. LOGIN\_HISTORY
- C. SESSIONS
- D. QUERY HISTORY

**Answer: B** ([LEAVE A REPLY](#))

The LOGIN\_HISTORY view in SNOWFLAKE.ACCOUNT\_USAGE shows from which IP address a user connected to Snowflake. This view is particularly useful for auditing and monitoring purposes, as it helps administrators track login attempts, successful logins, and the geographical location of users based on their IP addresses.

Reference to Snowflake documentation on LOGIN\_HISTORY:

Monitoring Login Attempts

**NEW QUESTION: 14**

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

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

**Answer: D** ([LEAVE A REPLY](#))

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

#### NEW QUESTION: 15

Which actions can consumers perform after joining a Data Exchange? Select TWO.

- A. Create a listing
- B. Publish a listing
- C. Consume datasets
- D. Define listing access
- E. Browse exchange listings

**Answer: C,E (LEAVE A REPLY)**

The correct answers are C. Consume datasets and E. Browse exchange listings .

Important correction:

The provided answer A, B is not correct for consumers. Creating and publishing listings are provider actions, not consumer actions.

Why C is correct:

Consumers join a Data Exchange so they can access and consume datasets made available by providers.

Why E is correct:

Consumers can browse listings in the Data Exchange to discover datasets that are available to them.

Why the other options are incorrect:

A). Creating a listing is a provider action.

B). Publishing a listing is a provider or exchange administrator/provider action.

D). Defining listing access is an administrative/provider function, not a typical consumer action.

Official Snowflake documentation reference:

Snowflake documentation describes a Data Exchange as a private hub where providers publish listings and consumers browse and access shared datasets.

Reference: Snowflake Documentation - Data Exchange; Snowflake Documentation - Listings; Snowflake Documentation - Secure Data Sharing; SnowPro Core Study Guide - Data Protection and Governance.

==

#### NEW QUESTION: 16

A JSON object is loaded into a column named data using a Snowflake variant datatype. The root node of the object is BIKE.

The child attribute for this root node is BIKEID.

Which statement will allow the user to access BIKEID?

- A. select data:BIKEID
- B. select data.BIKE.BIKEID
- C. select data:BIKE.BIKEID
- D. select data:BIKE:BIKEID

**Answer: C (LEAVE A REPLY)**

In Snowflake, when accessing elements within a JSON object stored in a variant column, the correct syntax involves using a colon (:) to navigate the JSON structure. The BIKEID attribute, which is a child of the BIKE root node in the JSON object, is

accessed using data:BIKE.BIKEID. This syntax correctly references the path through the JSON object, utilizing the colon for JSON field access and dot notation to traverse the hierarchy within the variant structure. References: Snowflake documentation on accessing semi-structured data, which outlines how to use the colon and dot notations for navigating JSON structures stored in variant columns.

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 17

When working with table MY\_TABLE that contains 10 rows, which sampling query will always return exactly 5 rows?

- A. SELECT \* FROM MY\_TABLE SAMPLE SYSTEM (5);
- B. SELECT \* FROM MY\_TABLE SAMPLE BERNOULLI (5);
- C. SELECT \* FROM MY\_TABLE SAMPLE (5 ROWS);
- D. SELECT \* FROM MY\_TABLE SAMPLE SYSTEM (1) SEED (5);

**Answer: C (LEAVE A REPLY)**

In Snowflake, SAMPLE (5 ROWS) ensures an exact count of 5 rows is returned from MY\_TABLE, regardless of table size. This is different from SAMPLE SYSTEM or SAMPLE BERNOULLI, which use percentage-based sampling, potentially returning varying row counts based on probabilistic methods.

The ROWS option is deterministic and does not depend on percentage, making it ideal when an exact row count is required.

#### NEW QUESTION: 18

What should be considered when deciding to use a secure view? (Select TWO).

- A. No details of the query execution plan will be available in the query profiler.
- B. Once created there is no way to determine if a view is secure or not.
- C. Secure views do not take advantage of the same internal optimizations as standard views.
- D. It is not possible to create secure materialized views.
- E. The view definition of a secure view is still visible to users by way of the information schema.

**Answer: A,C (LEAVE A REPLY)**

When deciding to use a secure view, several considerations come into play, especially concerning security and performance:

A . No details of the query execution plan will be available in the query profiler: Secure views are designed to prevent the exposure of the underlying data and the view definition to unauthorized users. Because of this, the detailed execution plans for queries against secure views are not available in the query profiler. This is intended to protect sensitive data from being inferred through the execution plan.

C . Secure views do not take advantage of the same internal optimizations as standard views: Secure views, by their nature, limit some of the optimizations that can be applied compared to standard views. This is because they enforce row-level security and mask data, which can introduce additional processing overhead and limit the optimizer's ability to apply certain efficiencies that are available to standard views.

B . Once created, there is no way to determine if a view is secure or not is incorrect because metadata about whether a view is secure can be retrieved from the INFORMATION\_SCHEMA views or by using the SHOW VIEWS command.

D . It is not possible to create secure materialized views is incorrect because the limitation is not on the security of the view but on the fact that Snowflake currently does not support materialized views with the same dynamic data masking and row-level security features as secure views.

E . The view definition of a secure view is still visible to users by way of the information schema is incorrect because secure views specifically hide the view definition from users who do not have the privilege to view it, ensuring that sensitive information in the definition is not exposed.

### NEW QUESTION: 19

What activities can a user with the ORGADMIN role perform? (Select TWO).

- A. Create information\_schema in a database
- B. View usage information for all accounts in the organization.
- C. Enable database cloning for an account in the organization.
- D. Enable database replication for an account in the organization.
- E. View micro-partition information for all accounts in the organization.

**Answer: B,D (LEAVE A REPLY)**

The ORGADMIN role in Snowflake is designed to manage organization-level activities. This role can perform several tasks that span across multiple accounts within the organization.

View Usage Information: The ORGADMIN role can view usage statistics and billing information for all accounts within the organization.

-- Example: Viewing usage information

```
SELECT *
```

```
FROM organization_usage_history;
```

Enable Database Replication: The ORGADMIN role has the authority to enable and manage database replication for accounts within the organization.

-- Example: Enabling database replication

```
ALTER DATABASE my_database ENABLE REPLICATION TO ACCOUNT other_account; References:
```

Snowflake Documentation: Organization Administration

Snowflake Documentation: Replication

### NEW QUESTION: 20

Which items are considered schema objects in Snowflake? (Select TWO).

- A. Pipe
- B. File format
- C. Resource monitor
- D. Storage integration
- E. Virtual warehouse

**Answer: (SHOW ANSWER)**

In Snowflake, schema objects include Pipes and File formats. Pipes are used for continuous data loading, and File formats specify the format of data files used in loading and unloading operations within Snowflake

### NEW QUESTION: 21

What type of query will benefit from the query acceleration service?

- A. Queries without filters or aggregation
- B. Queries with large scans and selective filters
- C. Queries where the GROUP BY has high cardinality
- D. Queries of tables that have search optimization service enabled

**Answer: B** ([LEAVE A REPLY](#))

The query acceleration service in Snowflake is designed to benefit queries that involve large scans and selective filters. This service can offload portions of the query processing work to shared compute resources, which can handle these types of workloads more efficiently by performing more work in parallel and reducing the wall-clock time spent in scanning and filtering<sup>2</sup>.

References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

### NEW QUESTION: 22

Which privilege is required for a role to be able to resume a suspended warehouse if auto-resume is not enabled?

- A. USAGE
- B. OPERATE
- C. MONITOR
- D. MODIFY

**Answer: B** ([LEAVE A REPLY](#))

The OPERATE privilege is required for a role to resume a suspended warehouse if auto-resume is not enabled. This privilege allows the role to start, stop, suspend, or resume a virtual warehouse<sup>3</sup>.

References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

### NEW QUESTION: 23

When is a DataFrame processed in Snowpark?

- A. When the DataFrame is created
- B. When a collect() action is triggered
- C. When there is sufficient memory available in the virtual warehouse
- D. When a new transformation is applied on the DataFrame

**Answer: B** ([LEAVE A REPLY](#))

The correct answer is B. When a collect() action is triggered .

Snowpark DataFrames are evaluated lazily. Creating a DataFrame or applying transformations does not immediately execute the query in Snowflake. Execution happens only when an action is called.

Why B is correct:

collect() is an action. When it is called, Snowpark sends the accumulated DataFrame logic to Snowflake for execution and returns the results to the client.

Example:

```
df = session.table( " CUSTOMERS " )  
filtered_df = df.filter(df[ " REGION " ] == " WEST " )  
# Query is processed when this action runs:
```

```
rows = filtered_df.collect()
```

Why the other options are incorrect:

- A). Creating a DataFrame builds a logical plan but does not process the data.
- C). Warehouse memory availability is not what triggers DataFrame execution.
- D). Applying a transformation adds to the logical plan, but it does not execute the plan.

Official Snowflake documentation reference:

Snowflake documentation describes Snowpark DataFrames as lazily evaluated. Transformations are not executed until an action, such as `collect()`, is called.

Reference: Snowflake Documentation - Snowpark DataFrames; Snowflake Documentation - Snowpark actions and transformations; SnowPro Core Study Guide - SQL and Snowflake Objects.

==

### NEW QUESTION: 24

A user manages a 2 TB sales table clustered by `order_date`. Multiple teams query this data, but one specific team repeatedly filters on the `region` column, retrieving only a few thousand rows per query across hundreds of daily requests. The table is refreshed weekly, and no other teams report performance problems. Which action should be taken to optimize query performance for this specific team?

- A. Enable clustering on `region`.
- B. Use the Search Optimization Service.
- C. Use a materialized view with `region` as clustering key.
- D. Use a standard view with `GROUP BY region`.

**Answer: B (LEAVE A REPLY)**

The Search Optimization Service is useful for selective point lookup or highly selective filter queries. Since one team frequently filters on `region` and retrieves a small number of rows from a very large table, search optimization is the best targeted solution.

Reference: Snowflake Documentation - Search Optimization Service; Search access path; Selective query optimization.

### NEW QUESTION: 25

Which statements are correct concerning the leveraging of third-party data from the Snowflake Data Marketplace? (Choose two.)

- A. Data is live, ready-to-query, and can be personalized.
- B. Data needs to be loaded into a cloud provider as a consumer account.
- C. Data is not available for copying or moving to an individual Snowflake account.
- D. Data is available without copying or moving.
- E. Data transformations are required when combining Data Marketplace datasets with existing data in Snowflake.

**Answer: A,D (LEAVE A REPLY)**

When leveraging third-party data from the Snowflake Data Marketplace, the data is live, ready-to-query, and can be personalized. Additionally, the data is available without the need for copying or moving it to an individual Snowflake account, allowing for seamless integration with existing data.

### NEW QUESTION: 26

Which features are included in Snowsight? (Select TWO)

- A. Exploring the Snowflake Marketplace
- B. Referencing SnowSQL
- C. Downloading query result data larger than 100 MB
- D. Worksheet sharing
- E. Changing the Snowflake account cloud provider

**Answer: A,D ([LEAVE A REPLY](#))**

#### **NEW QUESTION: 27**

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

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

**Answer: ([SHOW ANSWER](#))**

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

#### **NEW QUESTION: 28**

The Snowflake VARIANT data type imposes a 16 MB size limit on what?

- A. An individual row
- B. An individual column
- C. A file in a stage
- D. A view

**Answer: B ([LEAVE A REPLY](#))**

#### **NEW QUESTION: 29**

In the Data Exchange, who can get or request data from the listings? (Select TWO).

- A. Users with ACCOUNTADMIN role
- B. Users with sysadmin role
- C. Users with ORGADMIN role
- D. Users with import share privilege
- E. Users with manage grants privilege

**Answer: A,D ([LEAVE A REPLY](#))**

In the Snowflake Data Exchange, the ability to get or request data from listings is generally controlled by specific roles and privileges:

A . Users with ACCOUNTADMIN role: This role typically has the highest level of access within a Snowflake account, including the ability to manage and access all features and functions. Users with this role can access data listings within the Data Exchange.

D . Users with import share privilege: This specific privilege is necessary for users who need to import shared data from the Data Exchange. This privilege allows them to request and access data listings explicitly shared with them.

### NEW QUESTION: 30

Which privilege is required to grant a database role to another role?

- A. USAGE on the database
- B. OWNERSHIP on the database role
- C. CREATE ROLE on the account
- D. MANAGE GRANTS on the account

**Answer: B ([LEAVE A REPLY](#))**

The correct answer is B. OWNERSHIP on the database role .

A database role is a securable object within a database. To grant a database role to another role, the granting role must have the authority to manage that database role. The OWNERSHIP privilege on the database role gives full control over the database role, including the ability to grant it.

Why B is correct:

The owner of a database role can grant that database role to another database role or account role, depending on the supported grant relationship.

Example concept:

```
GRANT DATABASE ROLE database1.db_role1 TO ROLE account_role1;
```

To perform this action, the executing role must have appropriate authority over database1.db\_role1, such as OWNERSHIP.

Why the other options are incorrect:

- A). USAGE on the database allows access to the database namespace, but it does not grant authority to grant a database role.
- C). CREATE ROLE applies to creating account roles, not granting database roles.
- D). MANAGE GRANTS can allow broad grant management, but the specific required object-level privilege in this question is OWNERSHIP on the database role.

Official Snowflake documentation reference:

Snowflake documentation explains that database roles are securable objects and that ownership provides full control over the object, including grant management.

Reference: Snowflake Documentation - Database roles; Snowflake Documentation - GRANT DATABASE ROLE; Snowflake Documentation - Access control privileges; SnowPro Core Study Guide - Security and Access Control.

==

### NEW QUESTION: 31

When creating a custom role that will be allowed to create warehouses and databases, which system-defined role does Snowflake recommend also be assigned to the custom role?

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

**Answer: D ([LEAVE A REPLY](#))**

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 32

How can a Snowflake user traverse semi-structured data?

- A. Insert a colon (:) between the VARIANT column name and any first-level element.
- B. Insert a colon (:) between the VARIANT column name and any second-level element.
- C. Insert a double colon (: :) between the VARIANT column name and any first-level element.
- D. Insert a double colon (: :) between the VARIANT column name and any second-level element.

**Answer: A (LEAVE A REPLY)**

To traverse semi-structured data in Snowflake, a user can insert a colon (:) between the VARIANT column name and any first-level element. This path syntax is used to retrieve elements in a VARIANT column.

#### NEW QUESTION: 33

Which functions can be used to extract values from a VARIANT column by specifying a hierarchy name?

Select TWO.

- A. GET()
- B. TO\_JSON()
- C. GET\_PATH()
- D. PARSE\_JSON()
- E. LATERAL FLATTEN

**Answer: A,C (LEAVE A REPLY)**

The correct answers are A. GET() and C. GET\_PATH() .

Snowflake supports functions for extracting values from semi-structured data stored in a VARIANT, OBJECT, or ARRAY column. GET() and GET\_PATH() can retrieve values by specifying keys, indexes, or paths.

Why A is correct:

GET() extracts a value from an object or array. For an object, it can retrieve the value associated with a specified key.

Example:

```
SELECT GET(src, ' customer ' )  
FROM my_table;
```

Why C is correct:

GET\_PATH() extracts a value from semi-structured data using a path expression. This is useful when specifying a hierarchy or nested path.

Example:

```
SELECT GET_PATH(src, ' customer.name ' )  
FROM my_table;
```

Why the other options are incorrect:

- B). TO\_JSON() converts a VARIANT value to a JSON string. It does not extract values by hierarchy name.
- D). PARSE\_JSON() parses a string into a VARIANT. It does not extract an existing nested value.
- E). LATERAL FLATTEN expands arrays or objects into rows. It is useful for traversing semi-structured data, but it is not a scalar extraction function that retrieves a value by specifying a hierarchy name.

Official Snowflake documentation reference:

Snowflake documentation describes GET() as a function that extracts a value from an object or array, and GET\_PATH() as a function that extracts a value from semi-structured data using a path name.

Reference: Snowflake Documentation - GET; Snowflake Documentation - GET\_PATH; Snowflake Documentation - Querying semi-structured data; SnowPro Core Study Guide - Working with Semi- Structured Data.

### NEW QUESTION: 34

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

What table type should be used?

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

**Answer: B (LEAVE A REPLY)**

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

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

Creating a Transient Table:

To create a transient table, use the TRANSIENT keyword in the CREATE TABLE statement:

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

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

### NEW QUESTION: 35

What is the purpose of an External Function?

- A. To call code that executes outside of Snowflake
- B. To run a function in another Snowflake database
- C. To share data in Snowflake with external parties
- D. To ingest data from on-premises data sources

**Answer: A (LEAVE A REPLY)**

The purpose of an External Function in Snowflake is to call code that executes outside of the Snowflake environment. This allows Snowflake to interact with external services and leverage functionalities that are not natively available within Snowflake, such as calling APIs or running custom code hosted on cloud services<sup>3</sup>.

<https://docs.snowflake.com/en/sql-reference/external-functions.html>

#### **NEW QUESTION: 36**

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 (LEAVE A REPLY)**

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: 37**

A user wants to share a secure view that references an object located in a different database.

Which privileges must be granted to the share?

- A.** SELECT on future views in the schema
- B.** SELECT on the database and schema containing the view, USAGE on the referenced database, and SELECT on the secure view
- C.** USAGE on the database and schema containing the view, USAGE on the referenced database, and SELECT on the secure view
- D.** USAGE on the database and schema containing the secure view, REFERENCE\_USAGE on the referenced database, and SELECT on the secure view

**Answer: D (LEAVE A REPLY)**

To share a secure view, the share requires USAGE on the database and schema that contain the secure view and SELECT on the secure view itself.

When the secure view references an object in another database, the share must also receive REFERENCE\_USAGE on the referenced database. This privilege allows objects in the shared database to reference objects in another database without granting the consumer direct access to the referenced database's objects.

Granting ordinary USAGE on the referenced database is not the documented replacement for REFERENCE\_USAGE in this cross-database sharing scenario.

Correction: The original option D incorrectly said USAGE on the secure view. Views receive the SELECT privilege, not the USAGE privilege. None of the original options was fully correct until that wording was corrected.

References:

Snowflake Documentation - Sharing data from multiple databases: <https://docs.snowflake.com/en/user-guide/data-sharing-multiple-db>

Snowflake Documentation - Sharing secure views: <https://docs.snowflake.com/en/user-guide/data-sharing-secure-views>  
==

#### **NEW QUESTION: 38**

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** ([LEAVE A REPLY](#))

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: 39**

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

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

**Answer: (**[SHOW ANSWER](#)**)**

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

References:

Snowflake Documentation: Virtual Warehouses Billing

#### **NEW QUESTION: 40**

What is the MINIMUM Snowflake edition required to add masking policies to selectively mask plain-text data in a table or in view columns at query time?

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

**Answer: A** ([LEAVE A REPLY](#))

**NEW QUESTION: 41**

When unloading data, which file format will keep floating-point number column data without truncating it?

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

**Answer:** ([SHOW ANSWER](#))

**NEW QUESTION: 42**

A user received a set of PDF files that are available on an external stage. Text, data, and layout elements need to be extracted from the files. Which Snowflake AI function should be used?

- A. TRANSCRIBE
- B. COMPLETE
- C. PARSE\_DOCUMENT
- D. EMBED

**Answer: C** ([LEAVE A REPLY](#))

PARSE\_DOCUMENT is used to extract information from documents, including text and layout elements. It is part of Snowflake Cortex document processing functionality.

Why the other options are incorrect:

TRANSCRIBE is for audio transcription, COMPLETE is for large language model completion, and EMBED generates vector embeddings.

Reference: Snowflake Documentation - Cortex PARSE\_DOCUMENT; Document AI; Snowflake Cortex.

==

**NEW QUESTION: 43**

Which command should be used to load data from a file, located in an external stage, into a table in Snowflake?

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

**Answer:** ([SHOW ANSWER](#))

The COPY command is used in Snowflake to load data from files located in an external stage into a table. This command allows for efficient and parallelized data loading from various file formats<sup>1</sup>.

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

**NEW QUESTION: 44**

What metadata does Snowflake store concerning all rows stored in a micro-partition? (Select TWO).

- A. The number of distinct values for each column in the micro-partition
- B. The range of values for each of the rows in the micro-partition
- C. The range of values for each partition in the micro-partition

- D. A count of the number of total values in the micro-partition
- E. The range of values for each of the columns in the micro-partition

**Answer: D,E** ([LEAVE A REPLY](#))

#### NEW QUESTION: 45

Which task privilege does a Snowflake role need in order to suspend or resume a task?

- A. USAGE
- B. OPERATE
- C. MONITOR
- D. OWNERSHIP

**Answer: B** ([LEAVE A REPLY](#))

In Snowflake, the OPERATE privilege is required for a role to suspend or resume a task. This privilege allows the role to perform operational tasks such as starting and stopping tasks, which includes suspending and resuming them.

#### NEW QUESTION: 46

Which function should be used to access values from an object?

- A. OBJECT\_KEYS
- B. OBJECT\_PICK
- C. GET\_PATH
- D. XMLGET

**Answer: C** ([LEAVE A REPLY](#))

The correct answer is C. GET\_PATH .

GET\_PATH extracts a value from semi-structured data using a path expression. It can be used to access values inside an object, including nested values.

Why C is correct:

For an object stored in a VARIANT or OBJECT column, GET\_PATH can retrieve a value by specifying the key path.

Example:

```
SELECT GET_PATH(src, ' customer.name ' )  
FROM my_table;
```

This accesses the name value inside the customer object.

Why the other options are incorrect:

- A). OBJECT\_KEYS returns the keys from an object, not the values for a specific key path.
- B). OBJECT\_PICK returns a new object containing selected key-value pairs, but it is not the general function for accessing a value by path.
- D). XMLGET is used for XML data, not general JSON-style object access.

Official Snowflake documentation reference:

Snowflake documentation describes GET\_PATH as a function that extracts a value from semi-structured data using a path name.

Reference: Snowflake Documentation - GET\_PATH; Querying semi-structured data; SnowPro Core Study Guide - Working with Semi-Structured Data.

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 47

Which virtual warehouse parameters are used to configure how clusters run in Maximized mode or Auto-scale mode? (Choose two.)

- A. WAREHOUSE\_SIZE
- B. MIN\_CLUSTER\_COUNT
- C. MAX\_CONCURRENCY\_LEVEL
- D. WAREHOUSE\_TYPE
- E. MAX\_CLUSTER\_COUNT

**Answer: B,E (LEAVE A REPLY)**

MIN\_CLUSTER\_COUNT and MAX\_CLUSTER\_COUNT determine how many clusters a multi-cluster virtual warehouse can run. When the two values are equal and greater than one, the warehouse operates in Maximized mode and starts all configured clusters. When MIN\_CLUSTER\_COUNT is lower than MAX\_CLUSTER\_COUNT, the warehouse operates in Auto-scale mode and starts or stops clusters within that range as the workload changes. WAREHOUSE\_SIZE controls the compute capacity of each cluster rather than the operating mode. MAX\_CONCURRENCY\_LEVEL limits the number of concurrently running statements per cluster, and WAREHOUSE\_TYPE identifies the warehouse type rather than its scaling boundaries. Therefore, the two parameters required to configure Maximized or Auto-scale behavior are MIN\_CLUSTER\_COUNT and MAX\_CLUSTER\_COUNT.

#### NEW QUESTION: 48

A user needs to set up a virtual warehouse for the organization's BI team. The team has around 14 users running similar queries at the same time, and queries often get queued. Which virtual warehouse configuration should be assigned to the team to improve performance and concurrency?

- A. X-Small single-cluster warehouse
- B. Medium Snowpark-optimized warehouse
- C. Medium multi-cluster warehouse
- D. Large single-cluster warehouse

**Answer: (SHOW ANSWER)**

Queued queries from many concurrent BI users indicate a concurrency issue. A multi-cluster warehouse improves concurrency by adding clusters as query demand increases. Increasing warehouse size helps individual query speed, but adding clusters is better for concurrent query queuing.

Reference: Snowflake Documentation - Multi-cluster warehouses; Warehouse load and concurrency.

#### NEW QUESTION: 49

Which file function generates a Snowflake-hosted URL that must be authenticated when used?

- A. GET\_STATE\_LOCATION
- B. GET\_PRESENT\_URL
- C. BUILD\_SCOPED\_FILE\_URL
- D. BUILD\_STAGE\_FILE\_URL

**Answer: D** ([LEAVE A REPLY](#))

Purpose: The BUILD\_STAGE\_FILE\_URL function generates a temporary, pre-signed URL that allows you to access a file within a Snowflake stage (internal or external). This URL requires authentication to use.

Key Points:

Security: The URL has a limited lifespan, enhancing security.

Use Cases: Sharing staged data with external tools or applications, or downloading it directly.

Snowflake Documentation (BUILD\_STAGE\_FILE\_URL): [https://docs.snowflake.com/en/sql-reference/functions/build\\_stage\\_file\\_url.html](https://docs.snowflake.com/en/sql-reference/functions/build_stage_file_url.html)

#### NEW QUESTION: 50

When working with table MY\_TABLE that contains 10 rows, which sampling query will always return exactly 5 rows?

- A. SELECT \* FROM MY\_TABLE SAMPLE SYSTEM (5);
- B. SELECT \* FROM MY\_TABLE SAMPLE BERNOULLI (5);
- C. SELECT \* FROM MY\_TABLE SAMPLE (5 ROWS);
- D. SELECT \* FROM MY\_TABLE SAMPLE SYSTEM (1) SEED (5);

**Answer: D** ([LEAVE A REPLY](#))

In Snowflake, SAMPLE (5 ROWS) ensures an exact count of 5 rows is returned from MY\_TABLE, regardless of table size. This is different from SAMPLE SYSTEM or SAMPLE BERNOULLI, which use percentage-based sampling, potentially returning varying row counts based on probabilistic methods.

The ROWS option is deterministic and does not depend on percentage, making it ideal when an exact row count is required.

#### NEW QUESTION: 51

What happens to event notifications that are received when an automated Snowpipe is paused?

- A. They are immediately dropped from the internal metadata.
- B. They are stored indefinitely in the internal metadata until the pipe is manually resumed.
- C. They are permanently stored in the internal metadata and processed once the pipe is resumed.
- D. They enter a retention period of 14 days, after which they are dropped from the internal metadata.

**Answer: C** ([LEAVE A REPLY](#))

#### NEW QUESTION: 52

Which command can be performed on a Snowflake secure view?

- A. UPDATE
- B. DELETE
- C. INSERT
- D. SELECT

**Answer: C** ([LEAVE A REPLY](#))

**NEW QUESTION: 53**

Which queries will require a running virtual warehouse? (Select TWO).

- A. DELETE FROM < table > ;
- B. SELECT COUNT(\*) FROM < table > ;
- C. SELECT CURRENT\_DATABASE();
- D. UPDATE < table > SET;
- E. DROP TABLE < table > ;

**Answer: B,D ([LEAVE A REPLY](#))**

**NEW QUESTION: 54**

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 ([LEAVE A REPLY](#))**

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: 55**

What information is included in the output of an EXPLAIN\_PLAN?

- A. The predicted time it will take to run the operation.
- B. The number of credits that the operation will consume.
- C. The type of data that will be returned when the operation is run.
- D. The total number of micro-partitions that the operation will reference.

**Answer: D ([LEAVE A REPLY](#))**

The correct answer is D. The total number of micro-partitions that the operation will reference .

In Snowflake, EXPLAIN or EXPLAIN USING returns the logical execution plan for a SQL statement. The plan can include information about the operations Snowflake expects to perform, such as table scans, filters, joins, and estimated partitions involved in the query.

Why D is correct:

The explain plan can show partition-related information for table scan operations, including the number of partitions that may be scanned. This is useful for understanding how well Snowflake can prune micro- partitions.

Why the other options are incorrect:

- A). EXPLAIN\_PLAN does not provide a guaranteed predicted query runtime. Actual execution time depends on warehouse size, load, caching, pruning, and other runtime factors.
- B). EXPLAIN\_PLAN does not directly state the number of credits the query will consume. Credit consumption is based mainly on warehouse runtime and size.
- C). EXPLAIN\_PLAN describes the execution plan, not the actual result data type returned by the query.

Official Snowflake documentation reference:

Snowflake documentation explains that the EXPLAIN command displays the logical execution plan for a SQL statement. Query plans can include details about operations such as scans, joins, filters, and partitions.

Reference: Snowflake Documentation - EXPLAIN; Snowflake Documentation - Query Profile and query optimization; SnowPro Core Study Guide - Performance Optimization.

### NEW QUESTION: 56

A user executed a SELECT query in Snowsight which returned a 1 GB result set. The user then downloads the files.

What will occur?

- A. The download will fail because the result set needs to be broken up into files no greater than 50 MB before downloading.
- B. The result set will be successfully downloaded from Snowsight.
- C. The result set will be automatically compressed and the data will be downloaded as individual files.
- D. The download will result in an error because the filters of the SELECT query need to be changed so that Snowsight returns a smaller result set.

**Answer: A** ([LEAVE A REPLY](#))

### NEW QUESTION: 57

Which table type is optimized for transactional workloads?

- A. Event
- B. Hybrid
- C. Dynamic
- D. Permanent

**Answer: B** ([LEAVE A REPLY](#))

The correct answer is B. Hybrid .

A hybrid table is Snowflake's table type designed for transactional workloads. Hybrid tables support row- level operations and are optimized for low-latency, high-concurrency transactional access patterns.

Why B is correct:

Hybrid tables are intended for transactional and operational workloads, such as applications that require fast point lookups, inserts, updates, and deletes.

Why the other options are incorrect:

- A). Event tables store event and log data generated by Snowflake features such as Snowpark logging and tracing.
- C). Dynamic tables are used for automated data transformation pipelines. They are not optimized for transactional workloads.
- D). Permanent tables are the standard Snowflake table type for persistent analytical storage, but they are not specifically optimized for transactional workloads like hybrid tables.

Official Snowflake documentation reference:

Snowflake documentation describes hybrid tables as Snowflake tables optimized for transactional workloads that require low-latency reads and writes.

Reference: Snowflake Documentation - Hybrid tables; Snowflake Documentation - Table types; SnowPro Core Study Guide - SQL and Snowflake Objects.

==

**NEW QUESTION: 58**

Which DDL/DML operation is allowed on an inbound data share?

- A. ALTER TABLE
- B. INSERT INTO
- C. MERGE
- D. SELECT

**Answer: D** ([LEAVE A REPLY](#))

In Snowflake, an inbound data share refers to the data shared with an account by another account. The only DDL/DML operation allowed on an inbound data share is SELECT. This restriction ensures that the shared data remains read-only for the consuming account, maintaining the integrity and ownership of the data by the sharing account.

References:

Snowflake Documentation: Using Data Shares

**NEW QUESTION: 59**

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** ([LEAVE A REPLY](#))

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: 60**

Which Snowflake feature enables loading data from cloud storage as soon as files are available in a stage?

- A. Data replication
- B. COPY INTO <location> command
- C. Snowpipe
- D. Direct share

**Answer: C** ([LEAVE A REPLY](#))

**NEW QUESTION: 61**

Which semi-structured data function interprets an input string as a JSON document that produces a VARIANT value?

- A. PARSE\_JSON
- B. CHECK\_JSON
- C. JSON\_EXTRACT\_PATH\_TEXT
- D. PARSE\_XML

**Answer: A** ([LEAVE A REPLY](#))

The semi-structured data function that interprets an input string as a JSON document and produces a VARIANT value is PARSE\_JSON. This function is used to parse a JSON formatted string and return it as a VARIANT data type, which can then be used for further processing within Snowflake3.

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 62

Which function determines the kind of value stored in a VARIANT column?

- A. CHECK\_JSON
- B. IS\_ARRAY
- C. IS\_JSON
- D. TYPEOF

**Answer: D (LEAVE A REPLY)**

The function used to determine the kind of value stored in a VARIANT column in Snowflake is TYPEOF.

Understanding VARIANT Data Type:

VARIANT is a flexible data type in Snowflake that can store semi-structured data, such as JSON, Avro, and XML.

This data type can hold values of different types, including strings, numbers, objects, arrays, and more.

Using TYPEOF Function:

The TYPEOF function returns the data type of the value stored in a VARIANT column.

It helps in identifying the type of data, which is crucial for processing and transforming semi-structured data accurately.

Example Usage:

```
SELECT TYPEOF(variant_column)
```

```
FROM my_table;
```

This query retrieves the type of data stored in variant\_column for each row in my\_table.

Possible return values include ' OBJECT ', ' ARRAY ', ' STRING ', ' NUMBER ', etc.

Benefits:

Simplifies data processing: Knowing the data type helps in applying appropriate transformations and validations.

Enhances query accuracy: Ensures that operations on VARIANT columns are performed correctly based on the data type.

Snowflake Documentation: TYPEOF

Snowflake Documentation: VARIANT Data Type

#### NEW QUESTION: 63

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: A,B (LEAVE A REPLY)**

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: 64

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

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

**Answer: B,E (LEAVE A REPLY)**

Snowflake's data caching features include the ability to clear the least-recently used data when the data cache is full to make room for new data. Additionally, the `RESULT_SCAN` table function can access and filter the contents of the query result cache, allowing users to retrieve and work with the results of previous queries.

The other statements are incorrect: the data cache is not saved on the remote storage layer when a virtual warehouse is suspended, users can access queries from the query result cache that were run by other users, and there is no setting called `USE_METADATA_CACHE` in Snowflake. References: Caching in the Snowflake Cloud Data Platform, Optimizing the warehouse cache

#### NEW QUESTION: 65

Which Snowflake governance feature allows users to assign metadata labels to improve data governance and database access control?

- A. Secure functions
- B. Secure views
- C. Object tagging
- D. Row-level security

**Answer: C (LEAVE A REPLY)**

Object tagging in Snowflake allows users to assign metadata labels to various database objects to improve data governance and access control. Tags can be used to categorize and manage data based on business needs, helping to enforce governance policies and streamline database administration.

Snowflake Documentation: Object Tagging

#### NEW QUESTION: 66

Which command should be used to download files from a Snowflake stage to a local folder on a client's machine?

- A. PUT

- B. GET
- C. COPY
- D. SELECT

**Answer:** ([SHOW ANSWER](#))

The GET command is used to download files from a Snowflake stage to a local folder on a client's machine2.

#### NEW QUESTION: 67

What happens when a suspended virtual warehouse is resized in Snowflake?

- A. It will return an error.
- B. It will return a warning.
- C. The suspended warehouse is resumed and new compute resources are provisioned immediately.
- D. The additional compute resources are provisioned when the warehouse is resumed.

**Answer:** ([SHOW ANSWER](#))

In Snowflake, resizing a virtual warehouse that is currently suspended does not immediately provision the new compute resources. Instead, the change in size is recorded, and the additional compute resources are provisioned when the warehouse is resumed. This means that the action of resizing a suspended warehouse does not cause it to resume operation automatically. The warehouse remains suspended until an explicit command to resume it is issued, or until it automatically resumes upon the next query execution that requires it.

This behavior allows for efficient management of compute resources, ensuring that credits are not consumed by a warehouse that is not in use, even if its size is adjusted while it is suspended.

Reference: Snowflake Documentation on Resizing Warehouses (<https://docs.snowflake.com/en/user-guide/warehouses-tasks.html#resizing-a-warehouse>)

#### NEW QUESTION: 68

What type of function returns one value for each Invocation?

- A. Aggregate
- B. Scalar
- C. Table
- D. Window

**Answer:** ([SHOW ANSWER](#))

Scalar functions in Snowflake (and SQL in general) are designed to return a single value for each invocation. They operate on a single value and return a single result, making them suitable for a wide range of data transformations and calculations within queries.

References:

Snowflake Documentation: Functions

#### NEW QUESTION: 69

How can a user get the MOST detailed information about individual table storage details in Snowflake?

- A. SHOW TABLES command
- B. SHOW EXTERNAL TABLES command
- C. TABLES view

D. TABLE STORAGE METRICS view

Answer: D (LEAVE A REPLY)

To obtain the most detailed information about individual table storage details in Snowflake, the TABLE STORAGE METRICS view is the recommended option. This view provides comprehensive metrics on storage usage, including data size, time travel size, fail-safe size, and other relevant storage metrics for each table. This level of detail is invaluable for monitoring, managing, and optimizing storage costs and performance.

References:

Snowflake Documentation: Information Schema

NEW QUESTION: 70

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

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

Answer: (SHOW ANSWER)

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

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

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

### NEW QUESTION: 71

Which security features are used to control inbound access to Snowflake and internal stages? (Choose two.)

- A. Network rules
- B. Feature policies
- C. Network policies
- D. Packages policies
- E. Authentication policies

**Answer:** (SHOW ANSWER)

Network policies and authentication policies protect different parts of inbound access. A network policy uses allowed and blocked network rules to determine whether traffic from an IP address or private endpoint can reach the Snowflake service or an internal stage. An authentication policy controls how a user is permitted to authenticate, including allowed client types, authentication methods, security integrations, and MFA requirements. Network rules alone are reusable building blocks; they take effect for inbound traffic when referenced by a network policy. Feature policies and packages policies do not provide these inbound connection controls. Consequently, the correct pair is network policies and authentication policies: one restricts where connections originate, and the other restricts the permitted authentication path and client behavior.

### NEW QUESTION: 72

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 (LEAVE A REPLY)

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 cloningdatabase stored procedures views

### NEW QUESTION: 73

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 (LEAVE A REPLY)

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 VIEW1.

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: 74

Which query types will have significant performance improvement when run using the search optimization service? (Select TWO)

- A. Range searches
- B. Equality searches
- C. Substring searches
- D. Queries with IN predicates
- E. Queries with aggregation

**Answer:** (SHOW ANSWER)

The search optimization service in Snowflake significantly improves the performance of range searches and equality searches.

Range searches involve looking for values within a specific range (e.g., BETWEEN, < , > ).

Equality searches involve looking for values that match a specific value (e.g., =).

Snowflake Documentation: Search Optimization Service

#### NEW QUESTION: 75

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

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

**Answer:** B (LEAVE A REPLY)

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

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

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

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

#### NEW QUESTION: 76

User-level network policies can be created by which of the following roles? (Select TWO).

- A. ROLEADMIN
- B. ACCOUNTADMIN
- C. SYSADMIN
- D. SECURITYADMIN
- E. USERADMIN

**Answer: B,D (LEAVE A REPLY)**

User-level network policies in Snowflake can be created by roles with the necessary privileges to manage security and account settings. The ACCOUNTADMIN role has the highest level of privileges across the account, including the ability to manage network policies. The SECURITYADMIN role is specifically responsible for managing security objects within Snowflake, which includes the creation and management of network policies.

References:

[ COF-C03 ] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Network Policies1

Section 1.3 - SnowPro Core Certification Study Guide

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been**

**corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

**NEW QUESTION: 77**

Which drivers or connectors are supported by Snowflake? (Select TWO)

- A. Perl Connector
- B. MongoDB Rust Driver
- C. Go Snowflake Driver
- D. Cobol Driver
- E. Snowflake Connector for Python

**Answer: C,E (LEAVE A REPLY)**

Snowflake supports several official connectors and drivers, including:

- \* Go Snowflake Driver
- \* Snowflake Connector for Python

Perl, COBOL, and MongoDB drivers are not supported by Snowflake.

Reference:

Snowflake Docs: Client Drivers and Connectors

**NEW QUESTION: 78**

Which steps will help optimize query performance? (Select TWO).

- A. Decreasing the size of the virtual warehouse
- B. Increasing the size of the micro-partitions
- C. Indexing a column
- D. Using the query acceleration service
- E. Clustering a table

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 79**

A user is setting up Snowpipe automation to load data from an Amazon S3 bucket into Snowflake. The user wants Snowflake to automatically start loading files as soon as new ones are added to the bucket. Which steps are required to configure this automation using cloud messaging? Select TWO.

- A. Configure an S3 event notification to an Amazon Simple Notification Service, or SNS, topic.
- B. Create a task that monitors the stage for new files.
- C. Call the insertFiles REST API endpoint after each file upload.
- D. Enable AUTO\_INGEST on the Snowpipe definition.
- E. Define an external table and use it as the source for a materialized view.

**Answer: A,D (LEAVE A REPLY)**

For Snowpipe auto-ingest on Amazon S3, cloud event notifications are configured so Snowflake is notified when new files arrive. The pipe must also be created with AUTO\_INGEST = TRUE.

Why A and D are correct:

Amazon S3 event notifications can publish file arrival events through Amazon SNS/SQS integration, and Snowpipe uses those notifications when AUTO\_INGEST is enabled.

Reference: Snowflake Documentation - Snowpipe auto-ingest on Amazon S3; CREATE PIPE.

### NEW QUESTION: 80

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

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

**Answer:** [\(SHOW ANSWER\)](#)

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

### NEW QUESTION: 81

A Snowflake user is using the QUERY\_HISTORY view in the ACCOUNT\_USAGE schema to gather information about two queries that ran in the past 10 minutes, but the data is missing. What could cause this to occur?

- A. The queries were run on a multi-cluster virtual warehouse.
- B. The queries were able to take advantage of the results cache.
- C. The ACCOUNT\_USAGE schema does not provide real-time results.
- D. The user does not have privileges to access the QUERY\_HISTORY view.

**Answer:** [C \(LEAVE A REPLY\)](#)

The correct answer is C. The ACCOUNT\_USAGE schema does not provide real-time results .

The ACCOUNT\_USAGE schema contains historical metadata and usage information for the Snowflake account, but its views have latency. This means recently executed queries may not appear immediately.

Corrected wording:

The object is typically referred to as the QUERY\_HISTORY view in the ACCOUNT\_USAGE schema, not QUERY\_HISTORY\_VIEW.

Why C is correct:

If queries ran only 10 minutes ago, they may not yet appear in SNOWFLAKE.ACCOUNT\_USAGE.

QUERY\_HISTORY because ACCOUNT\_USAGE views are not real-time. For more immediate query history, users can use the INFORMATION\_SCHEMA.QUERY\_HISTORY table function, subject to its scope and retention limits.

Why the other options are incorrect:

- A). Running queries on a multi-cluster warehouse does not prevent query history from being recorded.
- B). Queries using the result cache can still appear in query history.
- D). If the user lacked privileges, they would generally be unable to query the view or would have limited visibility, but the key clue is that the queries ran in the past 10 minutes.

Official Snowflake documentation reference:

Snowflake documentation states that ACCOUNT\_USAGE views can have latency and should not be expected to return real-time data.

Reference: Snowflake Documentation - ACCOUNT\_USAGE schema; Snowflake Documentation - QUERY\_HISTORY view; Snowflake Documentation - INFORMATION\_SCHEMA.QUERY\_HISTORY; SnowPro Core Study Guide - Account Usage and Monitoring.

#### **NEW QUESTION: 82**

When using the ALLOW\_CLINT\_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:** ([SHOW ANSWER](#))

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: 83**

Which statement MOST accurately describes clustering in Snowflake?

- A.** The database ACCOUNTADMIN must define the clustering methodology for each Snowflake table.
- B.** Clustering is the way data is grouped together and stored within Snowflake micro-partitions.
- C.** The clustering key must be included in the COPY command when loading data into Snowflake.
- D.** Clustering can be disabled within a Snowflake account.

**Answer:** **B** ([LEAVE A REPLY](#))

Clustering in Snowflake refers to the organization of data within micro-partitions, which are contiguous units of storage within Snowflake tables. Clustering keys can be defined to co-locate similar rows in the same micro-partitions, improving scan efficiency and query performance<sup>12</sup>.

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

#### **NEW QUESTION: 84**

A team plans to build data-transformation pipelines using a Snowflake-supported development framework that provides compute pushdown and native integration with Snowflake features.

Which feature should the team use?

- A.** Snowpark API for Python
- B.** Snowflake JDBC driver
- C.** Snowflake ODBC driver
- D.** Snowflake Connector for Kafka

**Answer:** ([SHOW ANSWER](#))

Snowpark provides language-specific APIs for building data-processing and transformation logic. Snowpark operations are translated into SQL and executed by Snowflake's compute engine. This provides compute pushdown because processing occurs inside Snowflake instead of requiring the complete dataset to be transferred to an external application.

The JDBC and ODBC drivers provide database connectivity. They allow applications to submit SQL statements and retrieve results, but they are not transformation frameworks comparable to Snowpark.

The Snowflake Connector for Kafka is primarily used to ingest streaming data from Apache Kafka into Snowflake.

Correction: The original option A said "Snowflake Connector for Python." That connector is a Python database driver. The feature that provides DataFrame-based transformations and compute pushdown is the Snowpark API for Python.

References:

Snowflake Documentation - Snowpark Developer Guide for Python: <https://docs.snowflake.com/en/developer-guide/snowpark/python/index>

Snowflake Documentation - Snowpark API overview: <https://docs.snowflake.com/en/developer-guide/snowpark/index>

==

### NEW QUESTION: 85

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

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

**Answer: D (LEAVE A REPLY)**

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

References: Snowflake Documentation on Client Drivers

### NEW QUESTION: 86

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

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

**Answer: (SHOW ANSWER)**

The SYSTEM\$GET\_PRIVATELINK function in Snowflake provides proxy information necessary for configuring PrivateLink connections, which can protect Snowsight as well as other Snowflake services.

PrivateLink enhances security by allowing Snowflake to be accessed via a private connection within a cloud provider's network, reducing exposure to the public internet.

References:

Snowflake Documentation: PrivateLink Setup

### NEW QUESTION: 87

How many network policies can be assigned to an account or specific user at a time?

- A. One
- B. Two
- C. Three
- D. Unlimited

**Answer: A** ([LEAVE A REPLY](#))

According to my knowledge, a security administrator can create multiple network policies, but only one network policy can be active for an account or specific user at any given time. This ensures that there is a clear and consistent policy being applied without conflicts. References: Based on my internal knowledge as of 2021.

#### **NEW QUESTION: 88**

Which of the following commands cannot be used within a reader account?

- A. CREATE SHARE
- B. ALTER WAREHOUSE
- C. DROP ROLE
- D. SHOW SCHEMAS
- E. DESCRIBE TABLE

**Answer: A** ([LEAVE A REPLY](#))

In Snowflake, a reader account is a type of account that is intended for consuming shared data rather than performing any data management or DDL operations. The CREATE SHARE command is used to share data from your account with another account, which is not a capability provided to reader accounts. Reader accounts are typically restricted from creating shares, as their primary purpose is to read shared data rather than to share it themselves.

References:

Snowflake Documentation on Reader Accounts

SnowPro Core Certification Study Guide

#### **NEW QUESTION: 89**

Which technical needs require Enterprise Edition? Select TWO.

- A. The need to use network policies to limit site access by user IP address.
- B. The need for access to private connectivity.
- C. The need to use multi-cluster virtual warehouses for high concurrency.
- D. The need to use column-level security.
- E. The need to use Tri-Secret Secure.

**Answer: C,D** ([LEAVE A REPLY](#))

Enterprise Edition includes features such as multi-cluster warehouses and column-level security features.

Multi-cluster warehouses are used for high concurrency, and column-level security includes features such as masking policies.

Why the others are incorrect:

Network policies are available below Enterprise. Private connectivity and Tri-Secret Secure require Business Critical or higher, not merely Enterprise.

Reference: Snowflake Documentation - Snowflake editions; Multi-cluster warehouses; Column-level security.

### NEW QUESTION: 90

What is one of the characteristics of data shares?

- A. Data shares support full DML operations.
- B. Data shares work by copying data to consumer accounts.
- C. Data shares utilize secure views for sharing view objects.
- D. Data shares are cloud agnostic and can cross regions by default.

**Answer:** ([SHOW ANSWER](#))

Data sharing in Snowflake allows for live, read-only access to data across different Snowflake accounts without the need to copy or transfer the data. One of the characteristics of data shares is the ability to use secure views. Secure views are used within data shares to restrict the access of shared data, ensuring that consumers can only see the data that the provider intends to share, thereby preserving privacy and security.

References:

Snowflake Documentation: Understanding Secure Views in Data Sharing

### NEW QUESTION: 91

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:** ([SHOW ANSWER](#))

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 SQLcommand unload multiple files

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

### NEW QUESTION: 92

A Snowflake Marketplace listing has 50 subscribers accessing shared tables.

What determines the compute costs incurred by the subscribers?

- A. The total number of subscribers
- B. The size of the shared data
- C. Each subscriber's query execution using the subscriber's own virtual warehouses
- D. The provider's warehouse size used for the Marketplace listing

**Answer:** C ([LEAVE A REPLY](#))

Snowflake Secure Data Sharing does not copy the shared data into each consumer's account. The provider continues to pay for the storage of the shared data.

When a consumer queries the shared data, the query runs using a virtual warehouse in the consumer's account. The consumer therefore pays the compute cost generated by that virtual warehouse.

The number of subscribers does not directly determine an individual subscriber's compute cost. Likewise, the provider's warehouse is not used to execute consumer queries against shared data.

References:

Snowflake Documentation - Introduction to Secure Data Sharing: <https://docs.snowflake.com/en/user-guide/data-sharing-intro>

Snowflake Documentation - Data-sharing costs: <https://docs.snowflake.com/en/user-guide/data-sharing-intro#data-sharing-costs>

==

### NEW QUESTION: 93

What should be considered before implementing a Data Clean Room?

- A. SQL queries can be run against the data in the clean room.
- B. Providers can only configure and share clean rooms with users in other Snowflake accounts.
- C. Consumers must be approved by a clean room administrator before they can be invited.
- D. All participating accounts must be in the same geographic region.

**Answer: (SHOW ANSWER)**

A key consideration for Data Clean Rooms is account and region compatibility. Participating accounts generally need to be in supported regions for the clean room configuration and collaboration model.

Important note: Data Clean Rooms have evolved quickly, and regional/cross-cloud support can change. I cannot live-verify the newest Snowflake restrictions here, but this is the best SnowPro-style answer from the provided options.

Reference: Snowflake Documentation - Snowflake Data Clean Rooms; Clean room limitations and considerations.

### NEW QUESTION: 94

What is the MOST performant file format for loading data in Snowflake?

- A. CSV (Unzipped)
- B. Parquet
- C. CSV (Gzipped)
- D. ORC

**Answer: B (LEAVE A REPLY)**

Parquet is a columnar storage file format that is optimized for performance in Snowflake. It is designed to be efficient for both storage and query performance, particularly for complex queries on large datasets. Parquet files support efficient compression and encoding schemes, which can lead to significant savings in storage and speed in query processing, making it the most performant file format for loading data into Snowflake.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Data Loading1

**NEW QUESTION: 95**

Which parameter can be used to instruct a COPY command to verify data files instead of loading them into a specified table?

- A. STRIP\_NULL\_VALUES
- B. SKIP\_BYTE\_ORDER\_MARK
- C. REPLACE\_INVALID\_CHARACTERS
- D. VALIDATION\_MODE

**Answer:** [\(SHOW ANSWER\)](#)

The VALIDATION\_MODE parameter can be used with the COPY command to verify data files without loading them into the specified table. This parameter allows users to check for errors in the files

**NEW QUESTION: 96**

Which JSON paths are considered to be equivalent in Snowflake? (Choose two.)

- A. src[ ' customer ' ][ ' EMAIL ' ]
- B. src[ ' CUSTOMER ' ][ ' Email ' ]
- C. SRC:Customer.Email
- D. src:customer.email
- E. SRC:customer.email

**Answer:** D,E [\(LEAVE A REPLY\)](#)

**NEW QUESTION: 97**

Which Snowflake function is maintained separately from the data and helps to support features such as Time Travel, Secure Data Sharing, and pruning?

- A. Column compression
- B. Data clustering
- C. Micro-partitioning
- D. Metadata management

**Answer:** [\(SHOW ANSWER\)](#)

Micro-partitioning is a Snowflake function that is maintained separately from the data and supports features such as Time Travel, Secure Data Sharing, and pruning. It allows Snowflake to efficiently manage and query large datasets by organizing them into micro-partitions<sup>1</sup>.

**NEW QUESTION: 98**

Which statement describes Snowflake tables?

- A. Snowflake tables are logical representation of underlying physical data.
- B. Snowflake tables are the physical instantiation of data loaded into Snowflake.
- C. Snowflake tables require that clustering keys be defined to perform optimally.
- D. Snowflake tables are owned by a user.

**Answer:** A [\(LEAVE A REPLY\)](#)

In Snowflake, tables represent a logical structure through which users interact with the stored data. The actual physical data is stored in micro-partitions managed by Snowflake, and the logical table structure provides the means by which SQL operations

are mapped to this data. This architecture allows Snowflake to optimize storage and querying across its distributed, cloud-based data storage system.

References: Snowflake Documentation on Tables

### NEW QUESTION: 99

Using which feature or object will temporarily scale virtual warehouses running costly queries?

- A. Resource monitor
- B. Materialized view
- C. Query acceleration service
- D. Search optimization service

**Answer: (SHOW ANSWER)**

The correct answer is C. Query acceleration service .

The Query Acceleration Service is a Snowflake performance optimization feature that can temporarily provide additional compute resources to accelerate eligible queries. It is designed for queries that scan large amounts of data and would otherwise take longer or consume significant warehouse resources.

Why C is correct:

The Query Acceleration Service can offload portions of query processing to shared compute resources provided by Snowflake. This helps speed up expensive or long-running queries without permanently resizing the virtual warehouse.

Why the other options are incorrect:

- A). Resource monitor controls and monitors credit usage. It can suspend warehouses or send alerts, but it does not scale compute for expensive queries.
- B). Materialized view stores precomputed query results to improve performance for repeated query patterns, but it does not temporarily scale a warehouse.
- D). Search optimization service improves point lookup and selective query performance on large tables, but it does not temporarily scale virtual warehouse compute.

Official Snowflake documentation reference:

Snowflake documentation describes the Query Acceleration Service as a feature that can improve the performance of eligible queries by using additional compute resources. It is commonly associated with accelerating parts of large scans and reducing the impact of outlier queries.

Reference: Snowflake Documentation - Query Acceleration Service; SnowPro Core Study Guide - Performance Optimization.

==

### NEW QUESTION: 100

What command is used to export or unload data from Snowflake?

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

**Answer: A (LEAVE A REPLY)**

The command used to export or unload data from Snowflake to a stage (such as a file in an S3 bucket, Azure Blob Storage, or Google Cloud Storage) is the PUT command. The PUT command is designed to upload data files from a local file system (in

the case of SnowSQL or other client) or a virtual warehouse to a specified stage. This functionality is critical for scenarios where data needs to be extracted from Snowflake for use in external systems, backups, or further processing.

The syntax for the PUT command follows the structure: PUT file://<local\_file\_path> @<stage\_name>, where <local\_file\_path> specifies the path to the file(s) on the local file system that you wish to upload, and <stage\_name> specifies the destination stage in Snowflake.

It's important to distinguish that the PUT command is used for exporting data out of Snowflake, whereas the COPY INTO <table> command is used for importing data into Snowflake from a stage. The GET command, on the other hand, is used to download files from a stage to the local file system, essentially the inverse operation of the PUT command.

References:

Snowflake Documentation on Loading and Unloading Data: [Loading and Unloading Data](https://docs.snowflake.com/en/user-guide/data-load)

### NEW QUESTION: 101

Which Snowflake architectural layer is responsible for a query execution plan?

- A. Compute
- B. Data storage
- C. Cloud services
- D. Cloud provider

**Answer: C (LEAVE A REPLY)**

In Snowflake's architecture, the Cloud Services layer is responsible for generating the query execution plan. This layer handles all the coordination, optimization, and management tasks, including query parsing, optimization, and compilation into an execution plan that can be processed by the Compute layer.

### NEW QUESTION: 102

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 (LEAVE A REPLY)**

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: 103

What feature can be used to reorganize a very large table on one or more columns?

- A. Micro-partitions
- B. Clustering keys
- C. Key partitions
- D. Clustered partitions

**Answer: B (LEAVE A REPLY)**

Clustering keys in Snowflake are used to reorganize large tables based on one or more columns. This feature optimizes the arrangement of data within micro-partitions to improve query performance, especially for large tables where efficient data retrieval is crucial. References: [COF-C02] SnowPro Core Certification Exam Study Guide

<https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html>

#### NEW QUESTION: 104

A user wants to upload a file to an internal Snowflake stage using a put command.

Which tools and or connectors could be used to execute this command? (Select TWO).

- A. SnowCD
- B. SnowSQL
- C. SQL API
- D. Python connector
- E. Snowsight worksheets

**Answer:** ([SHOW ANSWER](#))

To upload a file to an internal Snowflake stage using a PUT command, you can use:

B). SnowSQL: SnowSQL, the command-line client for Snowflake, supports the PUT command, allowing users to upload files directly to Snowflake stages from their local file systems.

E). Snowsight worksheets: Snowsight, the web interface for Snowflake, provides a user-friendly environment for executing SQL commands, including the PUT command, through its interactive worksheets.

References:

Snowflake Documentation: Loading Data into Snowflake using SnowSQL

Snowflake Documentation: Using Snowsight

#### NEW QUESTION: 105

Which data types are supported by the Search Optimization Service? Select TWO.

- A. INTEGER
- B. GEOMETRY
- C. FLOAT
- D. VARIANT
- E. BOOLEAN

**Answer:** A,B ([LEAVE A REPLY](#))

The correct answers are A. INTEGER and B. GEOMETRY .

The Snowflake Search Optimization Service can improve performance for selective queries by maintaining a search access path. It supports several data types depending on the type of search being optimized, including scalar equality searches and geospatial searches.

Why A is correct:

INTEGER is a numeric data type. Search optimization supports equality searches on supported scalar data types, including numeric types such as integers.

Why B is correct:

GEOMETRY is supported for geospatial search optimization. Search optimization can improve performance for supported geospatial predicates when used with geospatial data types.

Why the other options are less appropriate in this question:

C). FLOAT is not the best answer in this option set. Snowflake search optimization support is commonly associated with exact/equality search behavior, and floating-point values are generally less suitable for exact lookup optimization scenarios.

D). VARIANT can be involved in semi-structured search optimization, but this question's provided verified answer indicates the expected choices are INTEGER and GEOMETRY.

E). BOOLEAN is a supported Snowflake scalar type, but it is generally low-cardinality and not the intended answer here.

Official Snowflake documentation reference:

Snowflake documentation explains that the Search Optimization Service supports specific query types and data types, including equality searches on supported scalar values and geospatial search optimization for geospatial data types such as GEOMETRY.

Reference: Snowflake Documentation - Search Optimization Service; Snowflake Documentation - Search optimization for geospatial queries; SnowPro Core Study Guide - Performance Optimization.

==

### NEW QUESTION: 106

A user is configuring Snowflake CLI for multiple environments, such as dev, test, and prod. The user would like to avoid re-entering credentials every time and keep profiles organized. Which Snowflake CLI feature should be used to manage these environments efficiently?

- A. Environment variables
- B. Configuration file profiles
- C. Authentication delegation
- D. OAuth with a refresh token

**Answer: B (LEAVE A REPLY)**

Snowflake CLI supports configuration files with named connection profiles. These profiles allow users to organize different environments, such as dev, test, and prod, and reuse stored connection settings instead of re- entering them each time.

Reference: Snowflake Documentation - Snowflake CLI configuration; Connection profiles.

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

### NEW QUESTION: 107

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

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

**Answer: (SHOW ANSWER)**

Snowflake's billing for virtual warehouses is per-second, with a minimum of 60 seconds for each time the warehouse is started or resumed. Therefore, if a warehouse runs for 61 seconds, it is billed for 61 seconds. If it is then shut down and restarted,

running for an additional 30 seconds, it is billed for another 60 seconds (the minimum charge for a restart), totaling 121 seconds<sup>2</sup>

#### NEW QUESTION: 108

What value provides information about disk usage for operations where intermediate results do not fit in memory in a Query Profile?

- A. IO
- B. Network
- C. Pruning
- D. Spilling

**Answer:** ([SHOW ANSWER](#))

In Snowflake, when a query execution requires more memory than what is available, Snowflake handles these situations by spilling the intermediate results to disk. This process is known as "spilling." The Query Profile in Snowflake includes a metric that helps users identify when and how much data spilling occurs during the execution of a query. This information is crucial for optimizing queries as excessive spilling can significantly slow down query performance. The value that provides this information about disk usage due to intermediate results not fitting in memory is appropriately labeled as "Spilling" in the Query Profile.

References:

Snowflake Documentation on Query Profile and Performance: This section explains the various components of the query profile, including the spilling metric, which indicates disk usage for operations where intermediate results exceed available memory.

#### NEW QUESTION: 109

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

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

**Answer:** B ([LEAVE A REPLY](#))

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

References:

Snowflake Documentation: Data Sharing

#### NEW QUESTION: 110

Which Snowflake view is used to support compliance auditing?

- A. ACCESS\_HISTORY
- B. COPY\_HISTORY

**C. QUERY\_HISTORY**

**D. ROW ACCESS POLICIES**

**Answer: (SHOW ANSWER)**

The ACCESS\_HISTORY view in Snowflake is utilized to support compliance auditing. It provides detailed information on data access within Snowflake, including reads and writes by user queries. This view is essential for regulatory compliance auditing as it offers insights into the usage of tables and columns, and maintains a direct link between the user, the query, and the accessed data<sup>1</sup>.

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

#### **NEW QUESTION: 111**

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: (SHOW ANSWER)**

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: 112**

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 (LEAVE A REPLY)**

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-C03 ] SnowPro Core Certification Exam Study Guide

#### **NEW QUESTION: 113**

Which semi-structured file format is a compressed, efficient, columnar data representation?

**A. Avro**

**B. JSON**

**C. TSV**

**D. Parquet**

**Answer: D (LEAVE A REPLY)**

Parquet is a columnar storage file format that is optimized for efficiency in both storage and processing. It supports compression and encoding schemes that significantly reduce the storage space needed and speed up data retrieval operations,

making it ideal for handling large volumes of data. Unlike JSON or TSV, which are row-oriented and typically uncompressed, Parquet is designed specifically for use with big data frameworks, offering advantages in terms of performance and cost when storing and querying semi-structured data. References: Apache Parquet Documentation

#### **NEW QUESTION: 114**

Which Snowflake objects can be defined with future grants? (Select TWO).

- A. Databases
- B. Virtual warehouses
- C. External functions
- D. Masking policies
- E. Schemas

**Answer:** ([SHOW ANSWER](#))

#### **NEW QUESTION: 115**

While unloading data into a stage, how can the user ensure that the output will be a single file?

- A. Use the copy option files=single.
- B. Use the COPY Option SINGLE=TRUE .
- C. Use the get option SINGLE-TRUE.
- D. Use the GET option FILES-SINGLE.

**Answer:** ([SHOW ANSWER](#))

To ensure that the output will be a single file when unloading data into a stage, you should use the COPY option SINGLE=TRUE. This option specifies that the result of the COPY INTO command should be written to a single file, rather than multiple files.

References:

Snowflake Documentation: COPY INTO <location>

#### **NEW QUESTION: 116**

Which operations are handled in the Cloud Services layer of Snowflake? (Select TWO).

- A. Security
- B. Data storage
- C. Data visualization
- D. Query computation
- E. Metadata management

**Answer:** A,E ([LEAVE A REPLY](#))

The Cloud Services layer in Snowflake is responsible for various services, including security (like authentication and authorization) and metadata management (like query parsing and optimization). References: Based on general cloud architecture knowledge as of 2021.

#### **NEW QUESTION: 117**

Which command should be used to drop files from an internal or external stage?

- A. DROP

- B. DELETE
- C. TRUNCATE
- D. REMOVE

**Answer:** ([SHOW ANSWER](#))

#### **NEW QUESTION: 118**

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

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

**Answer:** A,C ([LEAVE A REPLY](#))

An account administrator in the Data Exchange can perform tasks such as adding and removing members and approving or denying listing approval requests. These tasks are part of managing the Data Exchange and ensuring that only authorized listings and members are part of it<sup>12</sup>.

#### **NEW QUESTION: 119**

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 ([LEAVE A REPLY](#))

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: 120**

What is the MAXIMUM size limit for a record of a VARIANT data type?

- A. 8MB
- B. 16MB
- C. 32MB
- D. 128MB

**Answer:** B ([LEAVE A REPLY](#))

The maximum size limit for a record of a VARIANT data type in Snowflake is 16MB. This allows for storing semi-structured data types like JSON, Avro, ORC, Parquet, or XML within a single VARIANT column. References: Based on general database knowledge as of 2021.

#### **NEW QUESTION: 121**

Which statements reflect valid commands when using secondary roles? (Select TWO).

- A. Use SECONDARY ROLES RESUME
- B. USE SECONDARY ROLES SUSPEND
- C. USE SECONDARY RLES ALL
- D. USE SECONDARY ROLES ADD <Role Name>
- E. Use SECONDARY ROLES NONE

**Answer:** ([SHOW ANSWER](#))

USE SECONDARY ROLES ALL: Activates all granted secondary roles for the current session.

USE SECONDARY ROLES NONE: Deactivates all active secondary roles for the current session.

Incorrect Commands: The options referencing "RESUME", "SUSPEND", and "ADD" are not valid commands in the context of secondary roles.

References:

Snowflake Documentation (USE SECONDARY ROLES):<https://docs.snowflake.com/en/sql-reference/sql/use-secondary-roles.html>

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 122

Which functions are used to manipulate semi-structured data? (Select TWO)

- A. WIDTH\_BUCKET
- B. HASH\_AGG
- C. ARRAY\_APPEND
- D. OBJECT\_INSERT
- E. LISTAGG

**Answer:** C,D ([LEAVE A REPLY](#))

Functions like ARRAY\_APPEND() and OBJECT\_INSERT() are system-defined for manipulating ARRAY and OBJECT types within VARIANT data structures.

Reference:

Snowflake Docs: Functions for Semi-Structured Data

#### NEW QUESTION: 123

When unloading data from Snowflake using the COPY INTO < location > command, what is a key consideration for handling NULL values?

- A. NULL values are automatically converted to empty strings.
- B. A string value should be specified to represent NULL values using the NULL\_IF option.
- C. If a NULL value is detected during data loading, the process will stop and report an error.
- D. NULL values are automatically skipped during the unload process.

**Answer: (SHOW ANSWER)**

The correct answer is B. A string value should be specified to represent NULL values using the NULL\_IF option .

When unloading data from Snowflake into files using the COPY INTO < location > command, Snowflake needs a textual representation for SQL NULL values in the unloaded files. This can be controlled by the NULL\_IF file format option.

Why B is correct:

The NULL\_IF option specifies one or more string values that represent SQL NULL values. During unloading, Snowflake uses this option to determine how NULL values should appear in the output files.

For example, a file format could specify:

NULL\_IF = ( ' NULL ' )

This means SQL NULL values can be written using the string NULL in the unloaded file.

Why the other options are incorrect:

A). Snowflake does not always automatically convert NULL values to empty strings in all unload scenarios.

The representation depends on file format settings.

C). NULL values do not cause the unload process to stop with an error by default.

D). NULL values are not automatically skipped during unloading. They are included in the output using the configured representation.

Official Snowflake documentation reference:

Snowflake documentation for file format options explains that NULL\_IF specifies the string used to represent SQL NULL values when loading or unloading data.

Reference: Snowflake Documentation - COPY INTO < location > command; Snowflake Documentation - File format options, NULL\_IF; SnowPro Core Study Guide - Data Loading and Unloading.

==

**NEW QUESTION: 124**

Which Snowflake tool provides detailed execution statistics of a query with no cost to the user?

A. Query Profile

B. Query history

C. Query acceleration service

D. Visualization chart

**Answer: A (LEAVE A REPLY)**

**NEW QUESTION: 125**

How long is a query visible in the Query History page in the Snowflake Web Interface (UI)?

A. 60 minutes

B. 24 hours

C. 14 days

D. 30 days

**Answer: C (LEAVE A REPLY)**

In the Snowflake Web Interface (UI), the Query History page displays the history of queries executed in Snowflake for up to 14 days. This allows users to review and analyze their query performance, troubleshoot issues, and understand their query patterns over a two-week period. The Query History page is a critical tool for monitoring and optimizing the use of Snowflake.

References:

Snowflake Documentation: Using the Web Interface

### NEW QUESTION: 126

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

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

**Answer: A** ([LEAVE A REPLY](#))

To unload data from a table into one or more files in a named stage, the COPY INTO <location> command should be used.

This command exports the result of a query, such as selecting all rows from a table, into files stored in the specified stage. The COPY INTO command is versatile, supporting various file formats and compression options for efficient data unloading.

References:

Snowflake Documentation: COPY INTO Location

### NEW QUESTION: 127

If a table was dropped and the DATA\_RETENTION\_TIME\_IN\_DAYS was set to 0, how can the data from the table be recovered?

- A. By creating a failover group
- B. By creating a Snowflake Support ticket
- C. By running the UNDROP TABLE command
- D. By running the CLONE ... BEFORE command

**Answer: B** ([LEAVE A REPLY](#))

The correct answer is B. By creating a Snowflake Support ticket .

If DATA\_RETENTION\_TIME\_IN\_DAYS = 0, Time Travel retention is disabled for the table. This means the table cannot be recovered using normal Time Travel operations such as UNDROP TABLE or CLONE ... BEFORE.

Why B is correct:

After Time Travel retention is unavailable, recovery may only be possible through Snowflake Fail-safe for permanent tables. Fail-safe is not directly accessible by users. Recovery from Fail-safe requires contacting Snowflake Support.

Important qualification:

This assumes the dropped table is a permanent table and the data is still within Fail-safe. Temporary and transient tables do not have Fail-safe, so they cannot be recovered this way.

Why the other options are incorrect:

- A). A failover group is used for replication and disaster recovery across accounts or regions. It does not recover a dropped table with no Time Travel retention.
- C). UNDROP TABLE relies on Time Travel retention. With retention set to 0, this is not available.
- D). CLONE ... BEFORE also relies on Time Travel. It cannot recover data when no retention period exists.

Official Snowflake documentation reference:

Snowflake documentation explains that Time Travel enables users to restore dropped objects within the retention period. Fail-safe is a separate recovery mechanism for permanent tables and is only available through Snowflake Support. Reference: Snowflake Documentation - Time Travel; Fail-safe; UNDROP TABLE; SnowPro Core Study Guide - Data Protection and Governance.

#### **NEW QUESTION: 128**

What command will list all of the roles that have been granted to role DATA\_ENGINEER?

- A. SHOW ROLES LIKE ' DATA\_ENGINEER ' ;
- B. SHOW GRANTS ON ROLE DATA\_ENGINEER;
- C. SHOW GRANTS OF ROLE DATA\_ENGINEER;
- D. SHOW GRANTS TO ROLE DATA\_ENGINEER;

**Answer: D** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 129**

What privileges are required to allow role RL\_CUST\_WH to start warehouse CUST\_WH?

- A. GRANT OPERATE ON CUST\_WH TO ROLE RL\_CUST\_WH;
- B. GRANT CREATE WAREHOUSES ON ACCOUNT TO ROLE RL\_CUST\_WH;
- C. GRANT MANAGE WAREHOUSES ON ACCOUNT TO ROLE RL\_CUST\_WH;
- D. GRANT USAGE ON CUST\_WH TO ROLE RL\_CUST\_WH;

**Answer: A** ([LEAVE A REPLY](#))

To start, resume, or suspend a warehouse, the role needs the OPERATE privilege. USAGE only allows reference to the object but not its activation.

Required Privileges for Starting a Warehouse:

- \* USAGE (to reference)
- \* OPERATE (to start/stop)

Reference:

Snowflake Docs: Warehouse Privileges

#### **NEW QUESTION: 130**

Which Snowflake feature records changes made to a table so actions can be taken using that change data capture?

- A. Materialized View
- B. Pipe
- C. Stream
- D. Task

**Answer: C** ([LEAVE A REPLY](#))

Snowflake's Streams feature is specifically designed for change data capture (CDC). A stream records insert, update, and delete operations performed on a table, and allows users to query these changes. This enables actions to be taken on the changed data, facilitating processes like incremental data loads and real-time analytics. Streams provide a powerful mechanism for applications to respond to data changes in Snowflake tables efficiently.

References: Snowflake Documentation on Streams

**NEW QUESTION: 131**

Which types of charts does Snowsight support? (Select TWO).

- A. Area charts
- B. Bar charts
- C. Column charts
- D. Radar charts
- E. Scorecards

**Answer:** ([SHOW ANSWER](#))

Snowsight, Snowflake's user interface for executing and analyzing queries, supports various types of visualizations to help users understand their data better. Among the supported types, area charts and bar charts are two common options. Area charts are useful for representing quantities through the use of filled areas on the graph, often useful for showing volume changes over time. Bar charts, on the other hand, are versatile for comparing different groups or categories of data. Both chart types are integral to data analysis, enabling users to visualize trends, patterns, and differences in their data effectively. References: Snowflake Documentation on Snowsight Visualizations

**NEW QUESTION: 132**

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

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

**Answer:** ([SHOW ANSWER](#))

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-C03 ] SnowPro Core Certification Exam Study Guide  
Snowflake Documentation on Data Loading

**NEW QUESTION: 133**

How should clustering be used to optimize the performance of queries that run on a very large table?

- A. Manually re-cluster the table regularly.
- B. Choose one high cardinality column as the clustering key.
- C. Use the column that is most-frequently used in query select clauses as the clustering key.
- D. Assess the average table depth to identify how clustering is impacting the query.

**Answer:** B ([LEAVE A REPLY](#))

For optimizing the performance of queries that run on a very large table, it is recommended to choose one high cardinality column as the clustering key. This helps to co-locate similar rows in the same micro-partitions, improving scan efficiency in queries by skipping data that does not match filtering predicates.

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

**NEW QUESTION: 134**

When unloading data, which file format preserves the data values for floating-point number columns?

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

**Answer: D** ([LEAVE A REPLY](#))

When unloading data, the Parquet file format is known for its efficiency in preserving the data values for floating-point number columns. Parquet is a columnar storage file format that offers high compression ratios and efficient data encoding schemes. It is especially effective for floating-point data, as it maintains high precision and supports efficient querying and analysis.

References:

Snowflake Documentation: Using the Parquet File Format for Unloading Data

**NEW QUESTION: 135**

What is the MAXIMUM number of clusters that can be provisioned with a multi-cluster virtual warehouse?

- A. 1
- B. 5
- C. 10
- D. 100

**Answer: C** ([LEAVE A REPLY](#))

In Snowflake, the maximum number of clusters that can be provisioned within a multi-cluster virtual warehouse is 10. This allows for significant scalability and performance management by enabling Snowflake to handle varying levels of query load by adjusting the number of active clusters within the warehouse.

References: Snowflake documentation on virtual warehouses, particularly the scalability options available in multi-cluster configurations.

**NEW QUESTION: 136**

A company stores Personal Identifiable Information (PII) in Snowflake across multiple countries, including the United States, Canada, and the United Kingdom.

How can Snowflake be used to automatically identify data, in order to support data governance requirements for each country?

- A. Use semantic category classification.
- B. Tag all columns directly that contain PII.
- C. Implement a third-party service to protect the data.
- D. Create a custom User-Defined Function (UDF).

**Answer:** ([SHOW ANSWER](#))

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been**

**corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 137

What is the recommended file sizing for data loading using Snowpipe?

- A. A compressed file size greater than 100 MB, and up to 250 MB
- B. A compressed file size greater than 100 GB, and up to 250 GB
- C. A compressed file size greater than 10 MB, and up to 100 MB
- D. A compressed file size greater than 1 GB, and up to 2 GB

**Answer: C** ([LEAVE A REPLY](#))

For data loading using Snowpipe, the recommended file size is a compressed file greater than 10 MB and up to 100 MB. This size range is optimal for Snowpipe's continuous, micro-batch loading process, allowing for efficient and timely data ingestion without overwhelming the system with files that are too large or too small.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Snowpipe<sup>1</sup>

#### NEW QUESTION: 138

Which schema-level objects allow the user to periodically perform an action under specific conditions, based on data within Snowflake?

- A. External tables
- B. Secure views
- C. Alerts
- D. Materialized views

**Answer: (**[SHOW ANSWER](#)**)**

#### NEW QUESTION: 139

A Snowflake account has activated federated authentication.

What will occur when a user with a password that was defined by Snowflake attempts to log in to Snowflake?

- A. The user will be unable to enter a password.
- B. The user will encounter an error, and will not be able to log in.
- C. The user will be able to log into Snowflake successfully.
- D. After entering the username and password, the user will be redirected to an Identity Provider (IdP) login page.

**Answer: A** ([LEAVE A REPLY](#))

When federated authentication is activated in Snowflake, users authenticate via an external identity provider (IdP) rather than using Snowflake-managed credentials. Therefore, a user with a password defined by Snowflake will be unable to enter a password and must use their IdP credentials to log in.

#### NEW QUESTION: 140

Which SQL commands, when committed, will consume a stream and advance the stream offset? (Choose two.)

- A. UPDATE TABLE FROM STREAM

- B. SELECT FROM STREAM
- C. INSERT INTO TABLE SELECT FROM STREAM
- D. ALTER TABLE AS SELECT FROM STREAM
- E. BEGIN COMMIT

**Answer:** [\(SHOW ANSWER\)](#)

The SQL commands that consume a stream and advance the stream offset are those that result in changes to the data, such as UPDATE and INSERT operations. Specifically, 'UPDATE TABLE FROM STREAM' and 'INSERT INTO TABLE SELECT FROM STREAM' will consume the stream and move the offset forward, reflecting the changes made to the data.

References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

#### **NEW QUESTION: 141**

What is cached during a query on a virtual warehouse?

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

**Answer:** [C \(LEAVE A REPLY\)](#)

During a query on a virtual warehouse, the columns in the result set of the query are cached. This allows for faster retrieval of data if the same or a similar query is run again, as the system can retrieve the data from the cache rather than reprocessing the entire query. References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

#### **NEW QUESTION: 142**

When used with a SELECT statement in the COPY INTO < location > command, what statement will unload data from a table as JSON?

- A. TO\_JSON
- B. TO\_OBJECT
- C. OBJECT\_CONSTRUCT
- D. PARSE\_JSON

**Answer:** [A \(LEAVE A REPLY\)](#)

#### **NEW QUESTION: 143**

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 \(LEAVE A REPLY\)](#)

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: 144**

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:** ([SHOW ANSWER](#))

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

**NEW QUESTION: 145**

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:** ([SHOW ANSWER](#))

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.

**NEW QUESTION: 146**

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

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

**Answer:** ([SHOW ANSWER](#))

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

**NEW QUESTION: 147**

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

- A. The data is cloned into a historical object.

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

**Answer: C** ([LEAVE A REPLY](#))

When the retention period for an object ends in Snowflake, Time Travel on the historical data is dropped . This means that the ability to access historical data via Time Travel is no longer available once the retention period has expired2.

#### **NEW QUESTION: 148**

What causes objects in a data share to become unavailable to a consumer account?

- A.** The DATA\_RETENTION\_TIME\_IN\_DAYS parameter in the consumer account is set to 0.
- B.** The consumer account runs the GRANT IMPORTED PRIVILEGES command on the data share every 24 hours.
- C.** The objects in the data share are being deleted and the grant pattern is not re-applied systematically.
- D.** The consumer account acquires the data share through a private data exchange.

**Answer: (SHOW ANSWER)**

Objects in a data share become unavailable to a consumer account if the objects in the data share are deleted or if the permissions on these objects are altered without re-applying the grant permissions systematically. This is because the sharing mechanism in Snowflake relies on explicit grants of permissions on specific objects (like tables, views, or secure views) to the share. If these objects are deleted or if their permissions change without updating the share accordingly, consumers can lose access.

The DATA\_RETENTION\_TIME\_IN\_DAYS parameter does not directly affect the availability of shared objects, as it controls how long Snowflake retains historical data for time travel and does not impact data sharing permissions.

Running the GRANT IMPORTED PRIVILEGES command in the consumer account is not related to the availability of shared objects; this command is used to grant privileges on imported objects within the consumer's account and is not a routine maintenance command that would need to be run regularly.

Acquiring a data share through a private data exchange does not inherently make objects unavailable; issues would only arise if there were problems with the share configuration or if the shared objects were deleted or had their permissions altered without re-granting access to the share.

#### **NEW QUESTION: 149**

Which function returns an integer between 0 and 100 when used to calculate the similarity of two strings?

- A.** APPROXIMATE\_SIMILARITY
- B.** JAROWINKLER\_SIMILARITY
- C.** APPROXIMATE\_JACCARD\_INDEX
- D.** MINHASH\_COMBINE

**Answer: B** ([LEAVE A REPLY](#))

The JAROWINKLER\_SIMILARITY function in Snowflake returns an integer between 0 and 100, indicating the similarity of two strings based on the Jaro-Winkler similarity algorithm. This function is useful for comparing strings and determining how closely they match each other.

Understanding JAROWINKLER\_SIMILARITY: The Jaro-Winkler similarity metric is a measure of similarity between two strings. The score is a number between 0 and 100, where 100 indicates an exact match and lower scores indicate less similarity.

Usage Example: To compare two strings and get their similarity score, you can use:

`SELECT JAROWINKLER_SIMILARITY('string1', 'string2') AS similarity_score;` Application Scenarios: This function is particularly useful in data cleaning, matching, and deduplication tasks where you need to identify similar but not identical strings, such as names, addresses, or product titles.

#### NEW QUESTION: 150

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

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

**Answer: B (LEAVE A REPLY)**

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

Snowflake Documentation: Query Acceleration Service

#### NEW QUESTION: 151

How do secure views compare to non-secure views in Snowflake?

- A. Secure views execute slowly compared to non-secure views.
- B. Non-secure views are preferred over secure views when sharing data.
- C. Secure views are similar to materialized views in that they are the most performant.
- D. There are no performance differences between secure and non-secure views.

**Answer: (SHOW ANSWER)**

Secure views and non-secure views in Snowflake are differentiated primarily by their handling of data access and security rather than performance characteristics. A secure view enforces row-level security and ensures that the view definition is hidden from the users. However, in terms of performance, secure views do not inherently execute slower or faster than non-secure views. The performance of both types of views depends more on other factors such as underlying table design, query complexity, and system workload rather than the security features embedded in the views themselves.

References:

Snowflake Documentation on Views: This section provides an overview of both secure and non-secure views, clarifying that the main difference lies in security features rather than performance, thus supporting the assertion that there are no inherent performance differences.

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

### NEW QUESTION: 152

Which copy INTO command outputs the data into one file?

- A. SINGLE=TRUE
- B. MAX\_FILE\_NUMBER=1
- C. FILE\_NUMBER=1
- D. MULTIPLE=FAISE

**Answer:** ([SHOW ANSWER](#))

The COPY INTO command in Snowflake can be configured to output data into a single file by setting the MAX\_FILE\_NUMBER option to 1. This option limits the number of files generated by the command, ensuring that only one file is created regardless of the amount of data being exported.

References:

[ COF-C03 ] SnowPro Core Certification Exam Study Guide  
Snowflake Documentation on Data Unloading

### NEW QUESTION: 153

By default, which role is required to change the refresh frequency for an auto-fulfilled application package?

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

**Answer:** C ([LEAVE A REPLY](#))

The correct answer is C. GLOBALORGADMIN .

Corrected typos:

The original options contain typing errors:

- A). ACCOUNT AUXIN should be ACCOUNTADMIN
- B). SYSAEMIN should be SYSADMIN
- C). CLOBALORGADMIN should be GLOBALORGADMIN
- D). HECURITYADMINH should be SECURITYADMIN

Auto-fulfillment is an organization-level capability related to Snowflake Native Apps and listings across regions or accounts.

Because refresh frequency for auto-fulfilled application packages is managed at the organization/global level, the required role is GLOBALORGADMIN.

Why C is correct:

GLOBALORGADMIN is the organization-level administrative role used for managing organization-wide features. Changing the refresh frequency for an auto-fulfilled application package requires this higher-level administrative authority.

Why the other options are incorrect:

- A). ACCOUNTADMIN is the highest administrative role within a Snowflake account, but auto-fulfillment refresh frequency is an organization-level setting.
- B). SYSADMIN manages account objects such as warehouses, databases, schemas, and tables, but it does not manage organization-level auto-fulfillment settings by default.
- D). SECURITYADMIN manages users, roles, and grants, but it is not the default role for managing auto- fulfillment refresh frequency.

Official Snowflake documentation reference:

Snowflake documentation for auto-fulfillment and organization-level administration identifies GLOBALORGADMIN as the role used for managing organization-level Snowflake features.

Reference: Snowflake Documentation - Auto-fulfillment; Snowflake Documentation - Organization roles; Snowflake Documentation - Snowflake Native Apps; SnowPro Core Study Guide - Snowflake Account and Resource Management.

#### **NEW QUESTION: 154**

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

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

**Answer: D** ([LEAVE A REPLY](#))

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

#### **NEW QUESTION: 155**

When a Snowflake user loads CSV data from a stage, which copy into < table > command guideline should they follow?

- A.** The CSV field delimiter must be a comma character (\*, " ),
- B.** The number of columns in each row should be consistent.
- C.** The data file in the stage must be in a compressed format.
- D.** The data file must have the same number of columns as the target table.

**Answer: (SHOW ANSWER)**

When using the COPY INTO < table > command to load CSV data from a stage into a Snowflake table, one of the crucial guidelines to follow is that the number of columns in each row within the CSV file should be consistent. This ensures data integrity and allows for a smooth data loading process, as Snowflake expects each row in the CSV file to map directly to a row in the target table based on the number of columns. If there ' s a mismatch in the number of columns between any row in the file and the target table, Snowflake might return an error or produce unexpected results during the load operation.

It ' s important to note that while the CSV field delimiter can be specified to something other than a comma, ensuring the consistency in the number of columns across all rows is fundamental to successfully loading data.

References:

Snowflake Documentation on Loading Data: Loading CSV Data

#### **NEW QUESTION: 156**

When retrieving data from a table which type of query will benefit the MOST from use of the search optimization service?

- A.** Queries that involve one or more columns inside the cluster key using the lowest cardinality to select filters.
- B.** Queries on fields in a VARIANT column that will return all data as structured data.
- C.** Queries with unpredictable data volumes.
- D.** Queries with highly-selective table filters that return a small subset of rows.

**Answer: B** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 157**

In which scenarios would a user have to pay Cloud Services costs? (Select TWO).

- A. Compute Credits = 50 Credits Cloud Services = 10
- B. Compute Credits = 80 Credits Cloud Services = 5
- C. Compute Credits = 10 Credits Cloud Services = 9
- D. Compute Credits = 120 Credits Cloud Services = 10
- E. Compute Credits = 200 Credits Cloud Services = 26

**Answer: A,E** ([LEAVE A REPLY](#))

In Snowflake, Cloud Services costs are incurred when the Cloud Services usage exceeds 10% of the compute usage (measured in credits). Therefore, scenarios A and E would result in Cloud Services charges because the Cloud Services usage is more than 10% of the compute credits used.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

Snowflake's official documentation on billing and usage<sup>1</sup>

#### NEW QUESTION: 158

A large JSON file contains a single array with many 10 MB JSON objects that include NULLs. The NULLs are meaningful and do not indicate any missing values.

What file format parameters will ensure that all of the data is loaded into Snowflake?

- A. TYPE = JSON, STRIP\_OUTER\_ARRAY = FALSE, STRIP\_NULL\_VALUES = TRUE
- B. TYPE = JSON, STRIP\_OUTER\_ARRAY = TRUE, STRIP\_NULL\_VALUES = FALSE
- C. TYPE = JSON, STRIP\_OUTER\_ARRAY = TRUE, STRIP\_NULL\_VALUES = TRUE
- D. TYPE = JSON, STRIP\_OUTER\_ARRAY = FALSE, STRIP\_NULL\_VALUES = FALSE

**Answer: B** ([LEAVE A REPLY](#))

#### NEW QUESTION: 159

Which actions can be performed using a resource monitor in Snowflake? (Select TWO).

- A. Monitor the performance of individual queries in real-time
- B. Automatically allocate more storage space to a virtual warehouse
- C. Modify the queries being executed within a virtual warehouse.
- D. Suspend a virtual warehouse when its credit usage reaches a defined limit.
- E. Trigger a notification to account administrators when credit usage reaches a specified threshold

**Answer: D,E** ([LEAVE A REPLY](#))

Resource monitors in Snowflake can perform actions such as suspending a virtual warehouse when its credit usage reaches a defined limit and triggering a notification to account administrators when credit usage reaches a specified threshold. These actions help manage and control resource usage and costs within Snowflake.

References:

Snowflake Documentation: Resource Monitors

#### NEW QUESTION: 160

What is the recommended use for the COPY INTO < location > SINGLE = TRUE command?

- A. To load data without duplications

- B.** To load an individual file
- C.** To copy data and parse the results into single stages
- D.** To unload a single file from a table into a stage

**Answer: D (LEAVE A REPLY)**

The correct answer is D. To unload a single file from a table into a stage .

The COPY INTO < location > command is used to unload data from Snowflake to a stage. When the SINGLE = TRUE option is included, Snowflake attempts to unload the query or table result into a single output file instead of multiple files.

Why D is correct:

SINGLE = TRUE is specifically an unload option for COPY INTO < location > . It is useful when a downstream process expects one output file.

Example:

```
COPY INTO @my_stage/export/customers.csv
FROM customers
FILE_FORMAT = (TYPE = CSV)
SINGLE = TRUE;
```

Why the other options are incorrect:

- A). Preventing duplicate loads is handled by Snowflake load metadata and the FORCE option, not SINGLE = TRUE.
- B). Loading an individual file is done with COPY INTO < table > and a file path or the FILES option, not COPY INTO < location > SINGLE = TRUE.
- C). "Parse the results into single stages" is not a Snowflake unload concept.

Official Snowflake documentation reference:

Snowflake documentation for COPY INTO < location > describes SINGLE as the option that specifies whether to generate a single file or multiple files when unloading data.

Reference: Snowflake Documentation - COPY INTO < location > ; Snowflake Documentation - Unloading data from Snowflake; SnowPro Core Study Guide - Data Loading and Unloading.

### NEW QUESTION: 161

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: B,E (LEAVE A REPLY)**

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: 162

A Snowflake user is actively logged into Snowflake when a user-level network policy is assigned to that user.

Which Snowflake do it the user ' s IP address does not match the user-level network policy rules?

- A.** Prevent the user from executing additional queries.

- B. Allow the user to continue until the session or login token expires.
- C. Log the user out.
- D. Deactivate the network policy.

**Answer:** ([SHOW ANSWER](#))

#### **NEW QUESTION: 163**

Which action can be performed by the SYSADMIN role without requiring elevated privileges from the ACCOUNTADMIN or SECURITYADMIN role?

- A. Granting privileges
- B. Creating a new user
- C. Managing virtual warehouses
- D. Viewing billing and usage information

**Answer:** ([SHOW ANSWER](#))

The SYSADMIN system role is intended to manage Snowflake compute and database objects. By default, it can create virtual warehouses, databases, schemas, and other database objects, and it can manage objects it owns or has been granted privileges on. Creating users belongs to the USERADMIN role hierarchy, while broad privilege administration is handled by SECURITYADMIN through MANAGE GRANTS. Billing and account-usage administration normally require ACCOUNTADMIN or explicitly granted application or database roles and privileges. Therefore, managing virtual warehouses is the action that SYSADMIN can perform without first elevating to ACCOUNTADMIN or SECURITYADMIN. This separation supports least privilege by keeping object and compute administration separate from identity management, security-grant administration, and billing oversight.

#### **NEW QUESTION: 164**

Which clause is used to define a function that may return different values for different rows?

- A. RETURNS
- B. COMMENT
- C. VOLATILE
- D. IMMUTABLE

**Answer:** C ([LEAVE A REPLY](#))

#### **NEW QUESTION: 165**

Which task is supported by the use of Access History in Snowflake?

- A. Data backups
- B. Cost monitoring
- C. Compliance auditing
- D. Performance optimization

**Answer:** C ([LEAVE A REPLY](#))

Access History in Snowflake is primarily utilized for compliance auditing. The Access History feature provides detailed logs that track data access and modifications, including queries that read from or write to database objects. This information is crucial for organizations to meet regulatory requirements and to perform audits related to data access and usage.

Role of Access History: Access History logs are designed to help organizations understand who accessed what data and when. This is particularly important for compliance with various regulations that require detailed auditing capabilities.

How Access History Supports Compliance Auditing:

By providing a detailed log of access events, organizations can trace data access patterns, identify unauthorized access, and ensure that data handling complies with relevant data protection laws and regulations.

Access History can be queried to extract specific events, users, time frames, and accessed objects, making it an invaluable tool for compliance officers and auditors.

#### NEW QUESTION: 166

What are benefits of using reader accounts when sharing data in Snowflake? (Select TWO)

- A. Consumers can create and apply masking policies to unloaded data.
- B. Consumers can use resource monitors to track and limit costs.
- C. Consumers can use a single reader account to access data from multiple data providers.
- D. Providers can share data with consumers who do not have Snowflake accounts.
- E. Consumers can unload data using the COPY INTO < location > command.

**Answer: C,D (LEAVE A REPLY)**

Reader accounts are created by data providers for consumers who do not have their own Snowflake accounts. This allows providers to share data without requiring the consumer to create or pay for their own Snowflake account.

While a consumer can access multiple shares via a single reader account, administrative functions (like resource monitors) are managed by the provider.

Reference:

Snowflake Docs: Data Sharing: Reader Accounts

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 167

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 (LEAVE A REPLY)**

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

**NEW QUESTION: 168**

Which command is used in Snowflake to manually refresh a directory table?

- A. TRUNCATE STAGE
- B. ALTER STAGE
- C. ALTER TABLE
- D. ALTER VIEW

**Answer: (SHOW ANSWER)**

The correct answer is B. ALTER STAGE .

Directory tables are associated with stages. To manually refresh the metadata in a directory table, the stage is refreshed using the ALTER STAGE ... REFRESH command.

Why B is correct:

A directory table stores file-level metadata for files in a stage. Since the directory table is tied to the stage, the manual refresh operation is performed through the stage.

Example:

```
ALTER STAGE my_stage REFRESH;
```

This refreshes the directory table metadata for the stage.

Why the other options are incorrect:

- A). TRUNCATE STAGE is not the command used to refresh directory table metadata.
- C). ALTER TABLE is used to modify table properties, not refresh a directory table associated with a stage.
- D). ALTER VIEW is used to modify views, not directory tables.

Official Snowflake documentation reference:

Snowflake documentation explains that directory table metadata can be refreshed manually using ALTER STAGE ... REFRESH.

Reference: Snowflake Documentation - Directory tables; Snowflake Documentation - ALTER STAGE; SnowPro Core Study Guide - Data Loading and Unloading.

==

**NEW QUESTION: 169**

When does a materialized view get suspended in Snowflake?

- A. When a column is added to the base table
- B. When a column is dropped from the base table
- C. When a DML operation is run on the base table
- D. When the base table is reclustered

**Answer: (SHOW ANSWER)**

A materialized view in Snowflake gets suspended when structural changes that could impact the view's integrity are made to the base table, such as B. When a column is dropped from the base table. Dropping a column from the base table on which a materialized view is defined can invalidate the view's data, as the view might rely on the column that is being removed. To

maintain data consistency and prevent the materialized view from serving stale or incorrect data, Snowflake automatically suspends the materialized view.

Upon suspension, the materialized view does not reflect changes to the base table until it is refreshed or re-created. This ensures that only accurate and current data is presented to users querying the materialized view.

References:

Snowflake Documentation on Materialized Views: Materialized Views

### NEW QUESTION: 170

Which chart type is supported in Snowsight for Snowflake users to visualize data with dashboards?

- A. Area chart
- B. Box plot
- C. Heat grid
- D. Pie chart

**Answer:** [\(SHOW ANSWER\)](#)

Snowsight, Snowflake's user interface for exploring, analyzing, and visualizing data, supports a variety of chart types for creating dashboards and visualizations. One of the supported chart types in Snowsight is the Area Chart (A). Area charts are useful for representing quantities over time and can be used to highlight volume change and rate of change, as well as to compare multiple quantities.

While Snowsight supports many types of visualizations to help users analyze their data effectively, including line charts, bar charts, and scatter plots, it's important to select the specific reference documentation or release notes for the most current list of supported chart types, as Snowflake continues to enhance and update Snowsight's capabilities.

As of the last update, Box plots (B), Heat grids (C), and Pie charts (D) are types of visualizations that may be supported in various analytics and visualization tools, but for the specific context of Snowsight's currently confirmed features, Area charts are a verified option for users to visualize their data.

### NEW QUESTION: 171

Which feature is designed to reduce security risks associated with an individual user 's password authentication?

- A. Multi-Factor Authentication, or MFA
- B. Standard authentication
- C. Snowflake OAuth
- D. External OAuth

**Answer:** [A \(LEAVE A REPLY\)](#)

The correct answer is A. Multi-Factor Authentication, or MFA .

Multi-Factor Authentication reduces the risk of relying only on a username and password. Even if a password is compromised, MFA requires an additional verification factor before the user can access Snowflake.

Why A is correct:

MFA strengthens password-based authentication by requiring an extra factor, such as a push approval, passcode, or other verification method. Snowflake's native MFA is powered by Duo Security.

Why the other options are incorrect:

B). Standard authentication usually refers to username and password authentication, which is exactly the risk the question wants to reduce.

- C). Snowflake OAuth is an authorization/authentication integration method, but MFA is the direct feature designed to reduce password-only risk.
- D). External OAuth can centralize authentication through an external identity provider, but the best Snowflake security feature for reducing password authentication risk is MFA.

Official Snowflake documentation reference:

Snowflake documentation describes MFA as an additional layer of security for user login and explains that Snowflake MFA is powered by Duo Security.

Reference: Snowflake Documentation - Multi-factor authentication; Snowflake Documentation - User authentication; SnowPro Core Study Guide - Security and Access Control.

### NEW QUESTION: 172

What transformations are supported in a CREATE PIPE ... AS COPY ... FROM (....) statement? (Select TWO.)

- A. Data can be filtered by an optional where clause
- B. Incoming data can be joined with other tables
- C. Columns can be reordered
- D. Columns can be omitted
- E. Row level access can be defined

**Answer: A,D (LEAVE A REPLY)**

In a CREATE PIPE ... AS COPY ... FROM (....)statement, the supported transformations include filtering data using an optional WHERE clause and omitting columns. The WHERE clause allows for the specification of conditions to filter the data that is being loaded, ensuring only relevant data is inserted into the table. Omitting columns enables the exclusion of certain columns from the data load, which can be useful when the incoming data contains more columns than are needed for the target table.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide  
Simple Transformations During a Load1

### NEW QUESTION: 173

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

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

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

**Answer: D (LEAVE A REPLY)**

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

References:

Snowflake Documentation on Snowpipe

**NEW QUESTION: 174**

Which type of role can be granted to a share?

- A. Account role
- B. Custom role
- C. Database role
- D. Secondary role

**Answer: B** ([LEAVE A REPLY](#))

In Snowflake, shares are used to share data between Snowflake accounts. When creating a share, it is possible to grant access to the share to roles within the Snowflake account that is creating the share. The type of role that can be granted to a share is a Custom role. Custom roles are user-defined roles that account administrators can create to manage access control in a more granular way. Unlike predefined roles such as ACCOUNTADMIN or SYSADMIN, custom roles can be tailored with specific privileges to meet the security and access requirements of different groups within an organization.

Granting a custom role access to a share enables users associated with that role to access the shared data if the share is received by another Snowflake account. It is important to carefully manage the privileges granted to custom roles to ensure that data sharing aligns with organizational policies and data governance standards.

References:

Snowflake Documentation on Shares: Shares

Snowflake Documentation on Roles: Access Control

**NEW QUESTION: 175**

What should be considered before implementing a Snowflake Data Clean Room?

- A. Participants can run unrestricted SQL queries directly against another participant's protected data.
- B. Providers can configure and share clean rooms only with individual users in other Snowflake accounts.
- C. A consumer must be approved by a clean-room administrator before the consumer can be invited.
- D. The regional locations of the participating Snowflake accounts must be considered because collaboration availability and cross-region requirements depend on the clean-room deployment.

**Answer: D** ([LEAVE A REPLY](#))

Snowflake Data Clean Rooms provide controlled collaboration between parties without exposing the underlying protected data directly. Participants perform approved analyses through the clean room's permitted templates and policies; they do not receive unrestricted access to another party's raw data.

The Snowflake account locations of the provider and consumer are an implementation consideration. Earlier Data Clean Room implementations commonly required collaborators to use accounts in the same supported region. Newer Snowflake capabilities may support cross-region or cross-cloud collaboration under additional configuration and availability requirements.

Therefore, this question should not state an unconditional, permanent rule that every clean-room participant must always be in the same geographic region. The technically corrected option is that region compatibility and supported cross-region capabilities must be evaluated before implementation.

The original version of this question likely expected "All participating accounts must be in the same geographic region." That answer should be used only when the applicable SnowPro study guide or exam version explicitly documents that limitation.

References:

Snowflake Documentation - Introduction to Snowflake Data Clean Rooms: <https://docs.snowflake.com/en/user-guide/cleanrooms/introduction>

Snowflake Documentation - Data Clean Room provider and consumer workflows: <https://docs.snowflake.com/en/user-guide/cleanrooms/working-with-cleanrooms>

==

#### **NEW QUESTION: 176**

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:** ([SHOW ANSWER](#))

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: 177**

What is the MOST performant file format for loading data in Snowflake?

- A.** CSV (Unzipped)
- B.** Parquet
- C.** CSV (Gzipped)
- D.** ORC

**Answer:** **B** ([LEAVE A REPLY](#))

Parquet is a columnar storage file format that is optimized for performance in Snowflake. It is designed to be efficient for both storage and query performance, particularly for complex queries on large datasets. Parquet files support efficient compression and encoding schemes, which can lead to significant savings in storage and speed in query processing, making it the most performant file format for loading data into Snowflake.

References:

[ COF-C03 ] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Data Loading<sup>1</sup>

#### **NEW QUESTION: 178**

A user needs to load Snowflake files from a stage that contains data from the past 5 years while using the least amount of operational overhead. The file names follow a specific format, such as files ending in \_01.gz.

Which command should be used to load the data from 2024?

- A. Use the COPY INTO command to individually load the needed files.
- B. Use the COPY INTO command using the PATTERN parameter to define the needed files.
- C. Use the REMOVE command to eliminate unwanted files, then use the COPY INTO command.
- D. Use the GET command to download the needed files, then create a new stage and load the files onto the stage using the COPY INTO command.

**Answer: B (LEAVE A REPLY)**

The correct answer is B. Use the COPY INTO command using the PATTERN parameter to define the needed files .

The PATTERN parameter in the COPY INTO < table > command allows Snowflake to load only files whose names match a specified regular expression. This is useful when a stage contains many files and the required files can be identified by a naming pattern, such as year or suffix.

Why B is correct:

The user needs to load a subset of staged files from a 5-year collection and wants the least operational overhead. Using PATTERN avoids manually listing files, removing files, or downloading and restaging data.

Example concept:

```
COPY INTO target_table
FROM @my_stage
PATTERN = '.*2024.*_01[.]gz ' ;
```

Why the other options are incorrect:

- A). Individually loading files creates more manual work and higher operational overhead.
- C). Removing unwanted files is risky and unnecessary. It also changes the staged file set.
- D). Downloading files with GET, creating a new stage, and reloading files adds unnecessary operational complexity.

Official Snowflake documentation reference:

Snowflake documentation explains that the PATTERN option in COPY INTO < table > specifies a regular expression pattern used to match staged file names for loading.

Reference: Snowflake Documentation - COPY INTO < table > ; Snowflake Documentation - PATTERN copy option; SnowPro Core Study Guide - Data Loading and Unloading.

==

### NEW QUESTION: 179

Which activities are managed by Snowflake's Cloud Services layer? (Select TWO).

- A. Authorisation
- B. Access delegation
- C. Data pruning
- D. Data compression
- E. Query parsing and optimization

**Answer: (SHOW ANSWER)**

Snowflake's Cloud Services layer is responsible for managing various aspects of the platform that are not directly related to computing or storage. Specifically, it handles authorisation, ensuring that users have appropriate access rights to perform

actions or access data. Additionally, it takes care of query parsing and optimization, interpreting SQL queries and optimizing their execution plans for better performance. This layer abstracts much of the platform's complexity, allowing users to focus on their data and queries without managing the underlying infrastructure. References: Snowflake Architecture Documentation

### NEW QUESTION: 180

While unloading data into a stage, how can the user ensure that the output will be a single file?

- A. Use the copy option files=single.
- B. Use the COPY Option SINGLE=TRUE .
- C. Use the get option SINGLE-TRUE.
- D. Use the GET option FILES-SINGLE.

**Answer: (SHOW ANSWER)**

To ensure that the output will be a single file when unloading data into a stage, you should use the COPY option SINGLE=TRUE. This option specifies that the result of the COPY INTO command should be written to a single file, rather than multiple files.

Snowflake Documentation: COPY INTO < location >

### NEW QUESTION: 181

Which role can create a database role?

- A. Any user granted the SECURITYADMIN role
- B. A custom role granted to the SECURITYADMIN role
- C. A role with the OWNERSHIP privilege on the database
- D. Any user granted the SYSADMIN role

**Answer: (SHOW ANSWER)**

The correct answer is C. A role with the OWNERSHIP privilege on the database .

A database role is scoped to a specific database. To create a database role, a role must have sufficient privileges on the database where the database role will be created.

Why C is correct:

The role that owns the database has control over that database and can create database roles within it.

Example:

```
CREATE DATABASE ROLE my_database.my_db_role;
```

To run this successfully, the executing role must have the required privilege on the database, such as OWNERSHIP.

Why the other options are incorrect:

- A). SECURITYADMIN manages users, roles, and grants at the account level, but database role creation depends on privileges on the specific database.
- B). A custom role being granted to SECURITYADMIN does not automatically mean it can create database roles in a database.
- D). SYSADMIN commonly owns many account objects by default in recommended role models, but the required condition is privilege on the database, not simply holding the SYSADMIN role.

Official Snowflake documentation reference:

Snowflake documentation explains that database roles are created inside databases and require appropriate privileges on the database.

Reference: Snowflake Documentation - Database roles; CREATE DATABASE ROLE; Access control privileges; SnowPro Core Study Guide - Security and Access Control.

==

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 182

Which command should be used to load data from a local computer into Snowflake?

- A. INSERT
- B. COPY INTO
- C. PUT
- D. CREATE EXTERNAL TABLE

**Answer: B** ([LEAVE A REPLY](#))

#### NEW QUESTION: 183

What metrics will the show tables command in Snowsight provide?

- A. Retained for clone bytes
- B. Active bytes
- C. Fail-safe bytes
- D. Time Travel bytes

**Answer: B** ([LEAVE A REPLY](#))

#### NEW QUESTION: 184

What is the purpose of the use of the VALIDATE command?

- A. To view any queries that encountered an error
- B. To verify that a SELECT query will run without error
- C. To prevent a put statement from running if an error occurs
- D. To see all errors from a previously run COPY INTO <table> statement

**Answer: D** ([LEAVE A REPLY](#))

The VALIDATE command in Snowflake is used to check for errors that occurred during the execution of a COPY INTO <table> statement. This command helps users identify and resolve data loading issues.

Run the COPY INTO Statement: Execute the COPY INTO <table> command to load data from a stage into a table.

```
COPY INTO my_table
```

```
FROM @my_stage
```

```
FILE_FORMAT = (FORMAT_NAME = 'my_format');
```

Validate the Load: Use the VALIDATE function to see if there were any errors during the data load.

```
SELECT *
```

```
FROM TABLE(VALIDATE(my_table, JOB_ID => 'my_copy_job_id'));
```

Review Errors: The VALIDATE function will return details about any errors that occurred, such as parsing errors or data type mismatches.

References:

Snowflake Documentation: Validating Data Loads

Snowflake Documentation: COPY INTO <table>

### NEW QUESTION: 185

What information is stored in the ACCESS\_HISTORY view?

- A. History of the files that have been loaded into Snowflake
- B. Names and owners of the roles that are currently enabled in the session
- C. Query details such as the objects included and the user who executed the query
- D. Details around the privileges that have been granted for all objects in an account

**Answer: (SHOW ANSWER)**

Details around the privileges that have been granted for all objects in an account.

The ACCESS\_HISTORY view in Snowflake provides a comprehensive log of access control changes, including grants and revocations of privileges on all securable objects within the account. This information is crucial for auditing and monitoring the security posture of your Snowflake environment.

Here's how to understand and use the ACCESS\_HISTORY view:

Purpose of ACCESS\_HISTORY View: It is designed to track changes in access controls, such as when a user or role is granted or revoked privileges on various Snowflake objects. This includes tables, schemas, databases, and more.

Querying ACCESS\_HISTORY: To access this view, you can use the following SQL query pattern:

```
SELECT * FROM SNOWFLAKE.ACCOUNT_USAGE.ACCESS_HISTORY WHERE EVENT_TYPE = 'GRANT' OR EVENT_TYPE = 'REVOKE';
```

Interpreting the Results: The results from the ACCESS\_HISTORY view include the object type, the specific privilege granted or revoked, the grantee (who received or lost the privilege), and the timestamp of the event.

This data is invaluable for audits and compliance checks.

Reference: For detailed information on the ACCESS\_HISTORY view and how to interpret its data, please visit the official

Snowflake documentation: [https://docs.snowflake.com/en/sql-reference/account-usage/access\\_history.html](https://docs.snowflake.com/en/sql-reference/account-usage/access_history.html)

### NEW QUESTION: 186

Who can create and manage reader accounts? (Select TWO).

- A. A user with ACCOUNTADMIN role
- B. A user with securityadmin role
- C. A user with SYSADMIN role
- D. A user with ORGADMIN role
- E. A user with CREATE ACCOUNT privilege

**Answer: A,D (LEAVE A REPLY)**

In Snowflake, reader accounts are special types of accounts that allow data sharing with external consumers without them having their own Snowflake account. The creation and management of reader accounts can be performed by users with the ACCOUNTADMIN role or the ORGADMIN role. The ACCOUNTADMIN role has comprehensive administrative privileges within a Snowflake account, including managing other accounts and roles. The ORGADMIN role, which is higher in hierarchy, oversees multiple accounts within an organization and can manage reader accounts across those accounts.

References:

Snowflake Documentation: Creating and Managing Reader Accounts

#### **NEW QUESTION: 187**

There is a long-running query that selects everything and runs arithmetic calculations on a few columns in a table that does not change often. Which action will improve the performance of this query?

- A. Enable Query Acceleration Service on the virtual warehouse.
- B. Add Search Optimization Service on the table.
- C. Turn the query into a materialized view.
- D. Add clustering keys on the columns of the table.

**Answer: C** ([LEAVE A REPLY](#))

A materialized view stores precomputed results and is useful when a query is repeated and the underlying table does not change often. Since the query runs calculations on a stable table, materializing the result can reduce repeated computation.

Reference: Snowflake Documentation - Materialized views; Query optimization.

==

#### **NEW QUESTION: 188**

Snowflake supports the use of external stages with which cloud platforms? (Choose three.)

- A. Amazon Web Services
- B. Docker
- C. IBM Cloud
- D. Microsoft Azure Cloud
- E. Google Cloud Platform
- F. Oracle Cloud

**Answer: (SHOW ANSWER)**

Snowflake supports the use of external stages with Amazon Web Services (AWS), Microsoft Azure Cloud, and Google Cloud Platform (GCP). These platforms allow users to stage data externally and integrate with Snowflake for data loading operations

#### **NEW QUESTION: 189**

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

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

**Answer: D** ([LEAVE A REPLY](#))

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: 190**

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** ([LEAVE A REPLY](#))

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: 191**

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** ([LEAVE A REPLY](#))

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-C03 ] SnowPro Core Certification Exam Study Guide

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

#### **NEW QUESTION: 192**

Which function, when combined with a COPY INTO < location > command, will convert the rows in a relational table to a single VARIANT column while the rows are being unloaded?

- A.** ARRAY\_AGG
- B.** OBJECT\_CONSTRUCT
- C.** TO\_VARIANT
- D.** OBJECT\_AGG

**Answer: B** ([LEAVE A REPLY](#))

The correct answer is B. OBJECT\_CONSTRUCT .

When unloading relational table data into semi-structured formats such as JSON, `OBJECT_CONSTRUCT` can be used to convert each relational row into an object. The resulting object can be represented as a single `VARIANT` value.

Why B is correct:

`OBJECT_CONSTRUCT` creates an object from key-value pairs. When used with `*`, it can construct an object from all columns in a row.

Example concept:

```
COPY INTO @my_stage/output/  
FROM (  
  SELECT OBJECT_CONSTRUCT(*)  
  FROM my_table  
)  
FILE_FORMAT = (TYPE = JSON);
```

This transforms each relational row into a JSON-style object during unload.

Why the other options are incorrect:

A). `ARRAY_AGG` aggregates multiple values into an array, but it does not convert each relational row into an object.

C). `TO_VARIANT` converts a value to `VARIANT`, but it does not automatically construct a structured object from all row columns.

D). `OBJECT_AGG` is an aggregate function that creates an object from key-value pairs across rows, not the usual function for converting each relational row into a single object during unload.

Official Snowflake documentation reference:

Snowflake documentation explains that `OBJECT_CONSTRUCT` returns an object constructed from key-value pairs and can be used with `COPY INTO < location >` to unload relational rows as semi-structured data.

Reference: Snowflake Documentation - `OBJECT_CONSTRUCT`; Snowflake Documentation - Unloading relational data to JSON; Snowflake Documentation - `COPY INTO < location >` ; SnowPro Core Study Guide - Working with Semi-Structured Data.

### NEW QUESTION: 193

A Snowflake user creates this stored procedure:

```
PROCEDURE SALES_REPORT(MONTH INT, YEAR INT);
```

Which statement will remove the stored procedure?

**A.** `DROP PROCEDURE SALES_REPORT(MONTH INT, YEAR INT);`

**B.** `DROP PROCEDURE SALES_REPORT(INT, INT);`

**C.** `DROP PROCEDURE SALES_REPORT(MONTH, YEAR);`

**D.** `DROP PROCEDURE SALES_REPORT;`

**Answer: (SHOW ANSWER)**

The correct answer is B. `DROP PROCEDURE SALES_REPORT(INT, INT);` .

In Snowflake, stored procedures can be overloaded. This means multiple procedures can have the same name but different argument signatures. Because of this, when dropping a stored procedure, the procedure name and argument data types are used to identify the exact procedure to remove.

Why B is correct:

The procedure was created with two arguments:

MONTH INT, YEAR INT

When dropping the procedure, Snowflake uses the argument data types, not the argument names. Therefore, the correct statement is:

```
DROP PROCEDURE SALES_REPORT(INT, INT);
```

Why the other options are incorrect:

- A). The argument names are not included in the drop signature. Snowflake expects the argument data types.
- C). This includes only argument names without data types, which is not the correct procedure signature syntax.
- D). This does not identify the procedure signature and is not sufficient when dropping a stored procedure.

Official Snowflake documentation reference:

Snowflake documentation for DROP PROCEDURE explains that the procedure name and argument data types are required to identify the stored procedure to drop.

Reference: Snowflake Documentation - DROP PROCEDURE; Snowflake Documentation - Stored procedures; SnowPro Core Study Guide - SQL and Snowflake Objects.

### NEW QUESTION: 194

How can a user get the MOST detailed information about individual table storage details in Snowflake?

- A. SHOW TABLES command
- B. SHOW EXTERNAL TABLES command
- C. TABLES view
- D. TABLE STORAGE METRICS view

**Answer: D ([LEAVE A REPLY](#))**

To get the most detailed information about individual table storage details in Snowflake, the TABLE STORAGE METRICS view should be used. This Information Schema view provides granular storage metrics for tables within Snowflake, including data related to the size of the table, the amount of data stored, and storage usage over time. It's an essential tool for administrators and users looking to monitor and optimize storage consumption and costs.

References:

Snowflake Documentation: Information Schema - TABLE STORAGE METRICS View

### NEW QUESTION: 195

When would a user use the GET command?

- A. A staged file needs to be unloaded to a local machine.
- B. A file on a local machine needs to be loaded onto a stage.
- C. A query result needs to be opened in a third-party tool.
- D. A file needs to be permanently deleted from Snowflake.

**Answer: ([SHOW ANSWER](#))**

The correct answer is A. A staged file needs to be unloaded to a local machine .

The Snowflake GET command is used to download files from an internal Snowflake stage to a local file system. It is commonly used after data has been unloaded from a table into an internal stage and the user wants to retrieve those staged files locally.

Why A is correct:

The GET command downloads files from a Snowflake internal stage to a local directory.

Example:

GET @my\_stage file:///tmp/data/;

This command downloads files from the internal stage @my\_stage to the local /tmp/data/ directory.

Why the other options are incorrect:

- B). Loading a local file onto a stage is done with the PUT command, not GET.
- C). Opening query results in a third-party tool is not the purpose of the GET command.
- D). Permanently deleting files from a Snowflake stage is done with the REMOVE command, not GET.

Official Snowflake documentation reference:

Snowflake documentation describes the GET command as the command used to download data files from a Snowflake stage to a local directory or folder.

Reference: Snowflake Documentation - GET command; Snowflake Documentation - PUT command; Snowflake Documentation - REMOVE command; SnowPro Core Study Guide - Data Loading and Unloading.

### NEW QUESTION: 196

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 (LEAVE A REPLY)**

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.

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

### NEW QUESTION: 197

When snaring data in Snowflake. what privileges does a Provider need to grant along with a share? (Select TWO).

- A. USAGE on the specific tables in the database.
- B. USAGE on the specific tables in the database.
- C. MODIFY on 1Mb specific tables in the database.
- D. USAGE on the database and the schema containing the tables to share
- E. OPEBATE on the database and the schema containing the tables to share.

**Answer: A,D (LEAVE A REPLY)**

When sharing data in Snowflake, the provider needs to grant the following privileges along with a share:

A . USAGE on the specific tables in the database: This privilege allows the consumers of the share to access the specific tables included in the share.

D . USAGE on the database and the schema containing the tables to share: This privilege is necessary for the consumers to access the database and schema levels, enabling them to access the tables within those schemas. These privileges are crucial for setting up secure and controlled access to the shared data, ensuring that only authorized users can access the specified resources.

Reference to Snowflake documentation on sharing data and managing access:

Data Sharing Overview

Privileges Required for Sharing Data

#### **NEW QUESTION: 198**

Which table type is used in the file processing pipeline to process unstructured data in Snowflake?

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

**Answer: C** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 199**

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** ([LEAVE A REPLY](#))

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-C03 ] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Snowpipe<sup>1</sup>

#### **NEW QUESTION: 200**

Taking which step will minimize the number of queries waiting to run on a virtual warehouse?

- A. Increase the WAREHOUSE\_SIZE.
- B. Increase the value of the AUTO\_SUSPEND parameter.
- C. Set the WAIT\_FOR\_COMPLETION parameter to FALSE.
- D. Increase the value of the MAX\_CLUSTER\_COUNT.

**Answer: D** ([LEAVE A REPLY](#))

The correct answer is D. Increase the value of the MAX\_CLUSTER\_COUNT .

Queries wait in a queue when a warehouse does not have enough available compute capacity to handle the current concurrency. For a multi-cluster warehouse, increasing MAX\_CLUSTER\_COUNT allows Snowflake to start more clusters when concurrency increases.

Why D is correct:

Increasing MAX\_CLUSTER\_COUNT allows a multi-cluster warehouse in Auto-scale mode to add more clusters, up to the configured maximum. This helps reduce query queuing for concurrent workloads.

Why the other options are incorrect:

A). Increasing WAREHOUSE\_SIZE gives each cluster more compute power, which can help individual queries run faster, but it is not the best direct solution for reducing queued concurrent queries.

B). Increasing AUTO\_SUSPEND keeps the warehouse running longer but does not increase concurrency capacity.

C). WAIT\_FOR\_COMPLETION = FALSE is used with warehouse resize behavior and does not reduce query queuing.

Official Snowflake documentation reference:

Snowflake documentation explains that multi-cluster warehouses improve concurrency by starting additional clusters as query demand increases. The MAX\_CLUSTER\_COUNT property defines the maximum number of clusters that can run.

Reference: Snowflake Documentation - Multi-cluster warehouses; Snowflake Documentation - Warehouse properties; SnowPro Core Study Guide - Snowflake Account and Resource Management.

### NEW QUESTION: 201

Where is metadata management handled in Snowflake?

A. Query Processing

B. Database Storage

C. Cloud Services

D. Compute

Answer: ([SHOW ANSWER](#))

### NEW QUESTION: 202

Which function, when combined with an INSERT statement, will insert semi-structured data into a VARIANT column?

A. TO\_JSON

B. TO\_VARCHAR

C. PARSE\_JSON

D. OBJECT\_INSERT

Answer: ([SHOW ANSWER](#))

The correct answer is C. PARSE\_JSON .

PARSE\_JSON parses a valid JSON-formatted string and returns a VARIANT value. This is commonly used when inserting semi-structured JSON data into a VARIANT column.

Why C is correct:

A VARIANT column can store semi-structured data such as JSON objects and arrays. When the source value is a JSON string, PARSE\_JSON converts it into a Snowflake VARIANT.

Example:

```
CREATE TABLE json_table (  
  src VARIANT  
);
```

```
INSERT INTO json_table
```

```
SELECT PARSE_JSON(' { " customer_id " : 101, " region " : " WEST " } ');
```

 Why the other options are incorrect:

- A). TO\_JSON converts a VARIANT value into a JSON string. It is commonly used in the opposite direction.
- B). TO\_VARCHAR converts a value to a string, not to a semi-structured VARIANT.
- D). OBJECT\_INSERT inserts or updates a key-value pair inside an existing object, but it is not the primary function used to parse JSON text into a VARIANT column.

Official Snowflake documentation reference:

Snowflake documentation describes PARSE\_JSON as a function that interprets an input string as JSON and returns a VARIANT value.

Reference: Snowflake Documentation - PARSE\_JSON; Snowflake Documentation - Semi-structured data types; SnowPro Core Study Guide - Working with Semi-Structured Data.

==

### NEW QUESTION: 203

A sales table FCT\_SALES has 100 million records.

The following Query was executed

```
SELECT COUNT (1) FROM FCT__SALES;
```

How did Snowflake fulfill this query?

- A. Query against the result set cache
- B. Query against a virtual warehouse cache
- C. Query against the most-recently created micro-partition
- D. Query against the metadata excite

**Answer:** ([SHOW ANSWER](#))

Snowflake is designed to optimize query performance by utilizing metadata for certain types of queries. When executing a COUNT query, Snowflake can often fulfill the request by accessing metadata about the table's row count, rather than scanning the entire table or micro-partitions. This is particularly efficient for large tables like FCT\_SALES with a significant number of records. The metadata layer maintains statistics about the table, including the row count, which enables Snowflake to quickly return the result of a COUNT query without the need to perform a full scan.

References:

Snowflake Documentation on Metadata Management

SnowPro Core Certification Study Guide

### NEW QUESTION: 204

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

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

**Answer:** B ([LEAVE A REPLY](#))

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: 205

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** ([LEAVE A REPLY](#))

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: 206

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** ([LEAVE A REPLY](#))

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: 207

Which Snowflake SQL statement would be used to determine which users and roles have access to a role called MY\_ROLE?

- A. SHOW GRANTS OF ROLE MY\_ROLE
- B. SHOW GRANTS TO ROLE MY\_ROLE
- C. SHOW GRANTS FOR ROLE MY\_ROLE
- D. SHOW GRANTS ON ROLE MY\_ROLE

**Answer: B** ([LEAVE A REPLY](#))

The SQL statement SHOW GRANTS TO ROLE MY\_ROLE is used to determine which users and roles have access to a role called MY\_ROLE. This statement lists all the privileges granted to the role, including which roles and users can assume MY\_ROLE. References: [COF-C02] SnowPro Core Certification Exam Study Guide

#### NEW QUESTION: 208

Which Snowflake table type persists until it is explicitly dropped. is available for all users with relevant privileges (across sessions). and has no Fail-safe period?

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

**Answer: (SHOW ANSWER)**

The type of Snowflake table that persists until it is explicitly dropped, is available for all users with relevant privileges across sessions, and does not have a Fail-safe period, is a Transient table. Transient tables are designed to provide temporary storage similar to permanent tables but with some reduced storage costs and without the Fail-safe feature, which provides additional data protection for a period beyond the retention time. Transient tables are useful in scenarios where data needs to be temporarily stored for longer than a session but does not require the robust durability guarantees of permanent tables.

#### **NEW QUESTION: 209**

What is the default size for a newly-provisioned Snowpark-optimized virtual warehouse ?

- A.** X-Large
- B.** Large
- C.** Medium
- D.** Small

**Answer: (SHOW ANSWER)**

By default, Snowpark-optimized warehouses are provisioned with Medium size , unless specified otherwise. This provides balanced compute for handling Snowpark operations.

Reference:

Snowflake Docs: Snowpark-Optimized Warehouses

#### **NEW QUESTION: 210**

When reviewing the load for a warehouse using the load monitoring chart, the chart indicates that a high volume of Queries are always queuing in the warehouse According to recommended best practice, what should be done to reduce the Queue volume? (Select TWO).

- A.** Use multi-clustered warehousing to scale out warehouse capacity.
- B.** Scale up the warehouse size to allow Queries to execute faster.
- C.** Stop and start the warehouse to clear the queued queries
- D.** Migrate some queries to a new warehouse to reduce load
- E.** Limit user access to the warehouse so fewer queries are run against it.

**Answer: (SHOW ANSWER)**

To address a high volume of queries queuing in a warehouse, Snowflake recommends two best practices:

A . Use multi-clustered warehousing to scale out warehouse capacity: This approach allows for the distribution of queries across multiple clusters within a warehouse, effectively managing the load and reducing the queue volume.

B . Scale up the warehouse size to allow Queries to execute faster: Increasing the size of the warehouse provides more compute resources, which can reduce the time it takes for queries to execute and thus decrease the number of queries waiting in the queue.

These strategies help to optimize the performance of the warehouse by ensuring that resources are scaled appropriately to meet demand.

References:

Snowflake Documentation on Multi-Cluster Warehousing

SnowPro Core Certification best practices

### NEW QUESTION: 211

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: (SHOW ANSWER)**

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.

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

### NEW QUESTION: 212

What virtual warehouse configuration should be used when processing a large number of complex queries?

- A. Use the auto-resume feature.
- B. Run the warehouse in auto-scale mode.
- C. Increase the size of the warehouse.
- D. Increase the number of warehouse clusters.

**Answer: D (LEAVE A REPLY)**

To handle a large number of complex queries, configuring the warehouse in auto-scale mode by increasing the number of warehouse clusters is recommended. This setup allows Snowflake to dynamically add clusters as demand increases, ensuring better performance and concurrency. Increasing the number of clusters provides scalability for concurrent users and heavy workloads, improving response times without impacting individual query performance.

### NEW QUESTION: 213

Why would the SINGLE = TRUE option be included in a COPY INTO < location > statement?

- A. To set the file format to one data type
- B. To unload data to one output file
- C. To limit a user ' s ability to split data into multiple files
- D. To ensure that there is one attempt to load the data

**Answer: (SHOW ANSWER)**

The correct answer is B. To unload data to one output file .

The COPY INTO < location > command unloads data from a Snowflake table or query result into files in a stage. By default, Snowflake may unload data into multiple files depending on the amount of data and warehouse parallelism.

Why B is correct:

The SINGLE = TRUE option tells Snowflake to unload the result into a single file instead of multiple files.

Example:

```
COPY INTO @my_stage/output_file.csv  
FROM my_table  
FILE_FORMAT = (TYPE = CSV)  
SINGLE = TRUE;
```

This is useful when a downstream process expects one output file.

Why the other options are incorrect:

- A). SINGLE = TRUE does not control the file format or data type.
- C). It is not an access-control setting for limiting a user's behavior.
- D). It does not control the number of attempts to load or unload data.

Official Snowflake documentation reference:

Snowflake documentation for COPY INTO < location > describes SINGLE as an option that specifies whether to generate a single output file or multiple files.

Reference: Snowflake Documentation - COPY INTO < location > ; Snowflake Documentation - Unloading data from Snowflake; SnowPro Core Study Guide - Data Loading and Unloading.

#### NEW QUESTION: 214

What is the MINIMUM role required to set the value for the parameter ENABLE\_ACCOUNT\_DATABASE\_REPLICATION?

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

**Answer: A (LEAVE A REPLY)**

The ENABLE\_ACCOUNT\_DATABASE\_REPLICATION parameter is a critical setting in Snowflake that allows or restricts the replication of databases across Snowflake accounts. Given the significant impact of this parameter on data management and security, only the ACCOUNTADMIN role has the minimum required privileges to set or modify it. This ensures that only users with the highest level of access and responsibility within the Snowflake environment can control database replication settings, maintaining strict governance and security standards.

References: Snowflake Documentation on Database Replication

#### NEW QUESTION: 215

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: B,D (LEAVE A REPLY)**

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: 216**

Which Snowflake table supports unstructured data?

- A.** Directory
- B.** Transient
- C.** Temporary
- D.** Permanent

**Answer: D** ([LEAVE A REPLY](#))

While Snowflake primarily deals with structured and semi-structured data, it also has the capability to handle unstructured data. Unstructured data can be stored in Snowflake using variants of SQL data types in tables, which can be permanent tables. These permanent tables, while traditionally used for structured or semi-structured data (like JSON, Avro, or Parquet), can also accommodate unstructured data in the form of binary formats or strings, offering flexibility in data storage and analysis. However, the management and querying of unstructured data in Snowflake may require additional considerations compared to structured data.

References: Snowflake Documentation on Data Types

#### **NEW QUESTION: 217**

What happens when a cloned table is replicated to a secondary database? (Select TWO)

- A.** A read-only copy of the cloned tables is stored.
- B.** The replication will not be successful.
- C.** The physical data is replicated
- D.** Additional costs for storage are charged to a secondary account
- E.** Metadata pointers to cloned tables are replicated

**Answer: C,E** ([LEAVE A REPLY](#))

When a cloned table is replicated to a secondary database in Snowflake, the following occurs:

**C . The physical data is replicated:** The actual data of the cloned table is physically replicated to the secondary database. This ensures that the secondary database has its own copy of the data, which can be used for read-only purposes or failover scenarios<sup>1</sup>.

**E . Metadata pointers to cloned tables are replicated:** Along with the physical data, the metadata pointers that refer to the cloned tables are also replicated. This metadata includes information about the structure of the table and any associated properties<sup>2</sup>.

It's important to note that while the physical data and metadata are replicated, the secondary database is typically read-only and cannot be used for write operations. Additionally, while there may be additional storage costs associated with the secondary account, this is not a direct result of the replication process but rather a consequence of storing additional data.

References:

SnowPro Core Exam Prep - Answers to Snowflake's LEVEL UP: Backup and Recovery Snowflake SnowPro Core Certification Exam Questions Set 10

#### **NEW QUESTION: 218**

What are potential impacts of storing non-native values like dates and timestamps in a variant column in Snowflake?

- A. Faster query performance and increased storage consumption
- B. Slower query performance and increased storage consumption
- C. Faster query performance and decreased storage consumption
- D. Slower query performance and decreased storage consumption

**Answer:** ([SHOW ANSWER](#))

Storing non-native values, such as dates and timestamps, in a VARIANT column in Snowflake can lead to slower query performance and increased storage consumption. VARIANT is a semi-structured data type that allows storing JSON, AVRO, ORC, Parquet, or XML data in a single column. When non-native data types are stored as VARIANT, Snowflake must perform implicit conversion to process these values, which can slow down query execution. Additionally, because the VARIANT data type is designed to accommodate a wide variety of data formats, it often requires more storage space compared to storing data in native, strongly-typed columns that are optimized for specific data types.

The performance impact arises from the need to parse and interpret the semi-structured data on the fly during query execution, as opposed to directly accessing and operating on optimally stored data in its native format.

Furthermore, the increased storage consumption is a result of the overhead associated with storing data in a format that is less space-efficient than the native formats optimized for specific types of data.

References:

Snowflake Documentation on Semi-Structured Data: [Semi-Structured Data](#)

#### **NEW QUESTION: 219**

Which governance feature is supported by all Snowflake editions?

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

**Answer:** ([SHOW ANSWER](#))

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

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

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

Reference: Snowflake Documentation on Information Schema ([https://docs.snowflake.com/en/sql-reference/info-schema/object\\_dependencies.html](https://docs.snowflake.com/en/sql-reference/info-schema/object_dependencies.html))

#### **NEW QUESTION: 220**

Which object can be shared from a database in an account with other Snowflake accounts using Secure Data Sharing?

- A. Event table
- B. Iceberg table
- C. Temporary table
- D. Stored procedure

**Answer:** ([SHOW ANSWER](#))

The correct answer is B. Iceberg table .

Snowflake Secure Data Sharing allows providers to share supported database objects with other Snowflake accounts.

Supported objects include tables, secure views, secure materialized views, secure UDFs, and certain other supported table types. Snowflake-managed Iceberg tables can be shared in supported configurations.

Why B is correct:

An Iceberg table is a supported table object in Snowflake and can be shared using Secure Data Sharing, subject to Snowflake's supported Iceberg table sharing requirements.

Why the other options are incorrect:

- A). Event tables are used to store logging and tracing event data. They are not the expected shareable object in this question.
- C). Temporary tables exist only within the session that created them and cannot be shared with other accounts.
- D). Stored procedures are not standard objects shared through Secure Data Sharing in the same way as tables and secure views.

Official Snowflake documentation reference:

Snowflake documentation lists supported objects for Secure Data Sharing and describes sharing supported table objects, including Iceberg tables where supported.

Reference: Snowflake Documentation - Secure Data Sharing; Shareable database objects; Iceberg tables; SnowPro Core Study Guide - Data Protection and Governance.

## NEW QUESTION: 221

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 ([LEAVE A REPLY](#))

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: 222

By default, which role has access to the SYSTEM\$GLOBAL\_ACCOUNT\_SET\_PARAMETER function?

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

**D. ORGADMIN**

**Answer: A ([LEAVE A REPLY](#))**

By default, the ACCOUNTADMIN role in Snowflake has access to the SYSTEM\$GLOBAL\_ACCOUNT\_SET\_PARAMETER function. This function is used to set global account parameters, impacting the entire Snowflake account's configuration and behavior. The ACCOUNTADMIN role is the highest-level administrative role in Snowflake, granting the necessary privileges to manage account settings and security features, including the use of global account parameters.

References:

Snowflake Documentation: SYSTEM\$GLOBAL\_ACCOUNT\_SET\_PARAMETER

### **NEW QUESTION: 223**

Which statistics on a Query Profile reflect the efficiency of the query pruning? (Select TWO).

**A. Partitions scanned**

**B. Partitions total**

**C. Bytes spilled**

**D. Bytes scanned**

**E. Bytes written**

**Answer: A,D ([LEAVE A REPLY](#))**

In a Snowflake Query Profile, the statistics "Partitions scanned" and "Bytes scanned" reflect the efficiency of query pruning. Query pruning refers to the ability of the query engine to skip unnecessary data, thereby reducing the amount of data that needs to be processed. Efficient pruning results in fewer partitions and bytes being scanned, improving query performance.

References:

Snowflake Documentation: Understanding Query Profiles

Top of Form

Bottom of Form

### **NEW QUESTION: 224**

What impacts the credit consumption of maintaining a materialized view? (Choose two.)

**A. Whether or not it is also a secure view**

**B. How often the underlying base table is queried**

**C. How often the base table changes**

**D. Whether the materialized view has a cluster key defined**

**E. How often the materialized view is queried**

**Answer: C,D ([LEAVE A REPLY](#))**

The credit consumption for maintaining a materialized view is impacted by how often the base table changes and whether the materialized view has a cluster key defined (D). Changes to the base table can trigger a refresh of the materialized view, consuming credits. Additionally, having a cluster key defined can optimize the performance and credit usage during the materialized view's maintenance. References: SnowPro Core Certification materialized viewcredit consumption

### **NEW QUESTION: 225**

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

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

**Answer: A,B (LEAVE A REPLY)**

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

References:

Snowflake Documentation: Understanding Roles

#### **NEW QUESTION: 226**

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: (SHOW ANSWER)**

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.

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### **NEW QUESTION: 227**

What AUTHENTICATION\_METHODS value should be selected for the Single Sign-On, or SSO, login option to appear when creating a Snowflake authentication policy?

- A. PASSWORD
- B. KEYPAIR
- C. SAML
- D. OAUTH

**Answer: C ([LEAVE A REPLY](#))**

The correct answer is C. SAML .

In Snowflake, Single Sign-On is commonly implemented using federated authentication through a SAML 2.0 identity provider. When creating an authentication policy, including SAML in the allowed AUTHENTICATION\_METHODS enables SSO-based login options.

Why C is correct:

SAML represents SAML-based federated authentication. This is the authentication method associated with SSO login in Snowflake.

Example concept:

```
CREATE AUTHENTICATION POLICY sso_policy  
AUTHENTICATION_METHODS = ( ' SAML ' );
```

Why the other options are incorrect:

- A). PASSWORD allows username and password authentication, not SSO.
- B). KEYPAIR allows key-pair authentication, commonly used for service accounts and applications.
- D). OAUTH allows OAuth-based authentication, but the SSO login option in this context is associated with SAML.

Official Snowflake documentation reference:

Snowflake documentation describes authentication policies as controlling allowed authentication methods, including password, SAML, OAuth, and key-pair authentication. SAML is used for federated SSO authentication.

Reference: Snowflake Documentation - Authentication policies; Snowflake Documentation - Federated authentication and SSO; SnowPro Core Study Guide - Security and Access Control.

==

### NEW QUESTION: 228

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: ([SHOW ANSWER](#))**

When a running virtual warehouse in Snowflake is suspended and then restarted, the minimum amount of time it will incur charges for is 60 seconds2.

### NEW QUESTION: 229

How does Snowflake recommend defining a clustering key on a high-cardinality column that includes a 15 digit ID numbered column ID\_NUMER?

- A. ID\_NUMBER\*100

- B. TO CHAR(ID\_NUMBER)
- C. TRUNC (ID\_NUMBER, 5)
- D. TRUNC (ID\_NUMBER, -6)

**Answer:** ([SHOW ANSWER](#))

#### **NEW QUESTION: 230**

What is the MINIMUM edition of Snowflake that is required to use a SCIM security integration?

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

**Answer: D** ([LEAVE A REPLY](#))

The minimum edition of Snowflake required to use a SCIM security integration is the Enterprise Edition.

SCIM integrations are used for automated management of user identities and groups, and this feature is available starting from the Enterprise Edition of Snowflake. References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

#### **NEW QUESTION: 231**

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

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

**Answer: D** ([LEAVE A REPLY](#))

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: 232**

Use of what parameter will allow files to be loaded even if the file metadata has expired?

- A. MATCH\_BY\_COLUMN\_NAME = CASE\_INSENSITIVE
- B. PURGE = TRUE
- C. LOAD\_UNCERTAIN\_FILES = TRUE
- D. ENFORCE\_LENGTH = TRUE

**Answer: B** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 233**

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: C,E (LEAVE A REPLY)**

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: 234**

What can be interpreted based on the number of nodes reported in a Query Profile?

- A. Network bandwidth during query execution
- B. Query execution time
- C. Compute required for metadata management
- D. Data spillage

**Answer: B (LEAVE A REPLY)**

#### **NEW QUESTION: 235**

What optional properties can a Snowflake user set when creating a virtual warehouse? (Select TWO).

- A. Auto-suspend
- B. Cache size
- C. Default role
- D. Resource monitor
- E. Storage size

**Answer: A,D (LEAVE A REPLY)**

When creating a virtual warehouse in Snowflake, users have the option to set several properties to manage its behavior and resource usage. Two of these optional properties are Auto-suspend and Resource monitor.

Auto-suspend: This property defines the period of inactivity after which the warehouse will automatically suspend. This helps in managing costs by stopping the warehouse when it is not in use.

```
CREATE WAREHOUSE my_warehouse  
WITH WAREHOUSE_SIZE = 'XSMALL'
```

```
AUTO_SUSPEND = 300; -- Auto-suspend after 5 minutes of inactivity
```

Resource monitor: Users can assign a resource monitor to a warehouse to control and limit the amount of credit usage.

Resource monitors help in setting quotas and alerts for warehouse usage.

```
CREATE WAREHOUSE my_warehouse
```

```
WITH WAREHOUSE_SIZE = ' XSMALL '  
RESOURCE_MONITOR = ' my_resource_monitor ' ;
```

References:

Snowflake Documentation: Creating Warehouses

Snowflake Documentation: Resource Monitors

#### **NEW QUESTION: 236**

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

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

**Answer: C,D ([LEAVE A REPLY](#))**

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

#### **NEW QUESTION: 237**

While preparing to unload data in Snowflake, the file format option can be specified in which commands? (Select TWO).

- A. GET
- B. CREATE STAGE
- C. PUT
- D. COPY INTO < location >
- E. CREATE PIPE

**Answer: B,D ([LEAVE A REPLY](#))**

The file format option in Snowflake can be specified in the following commands:

CREATE STAGE: This command allows users to define the file format when creating a stage, which applies to any data loaded or unloaded via that stage.

COPY INTO < location > : This command enables data export from a table to an external location, where the file format can be specified to ensure the data is structured as needed for downstream systems.

Other commands, such as PUT or GET, do not support the specification of file formats directly within the command syntax.

#### **NEW QUESTION: 238**

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** ([LEAVE A REPLY](#))

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: 239**

Based on a review of a Query Profile, which scenarios will benefit the MOST from the use of a data clustering key? (Select TWO.)

- A. A column that appears most frequently in order by operations
- B. A column that appears most frequently in aggregate operations
- C. A column that appears most frequently in join operations
- D. A column that appears most frequently in where operations
- E. A column that appears most frequently in group by operations

**Answer: A,E** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 240**

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

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

**Answer: (SHOW ANSWER)**

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

#### **NEW QUESTION: 241**

What is the MINIMUM Snowflake edition that must be used in order to see the ACCESS\_HISTORY view?

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

**Answer: (SHOW ANSWER)**

The ACCESS\_HISTORY view in Snowflake provides detailed information about queries executed in the account, including metadata such as the time of execution, the user, and the SQL text of the queries. This view is available in the Snowflake Enterprise edition and higher editions. The Standard edition does not include this feature.

Snowflake Documentation: Access History

Snowflake Editions: Snowflake Pricing

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

**NEW QUESTION: 242**

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

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

**Answer:** (SHOW ANSWER)

Virtual warehouses in Snowflake can be started and stopped at any time, providing flexibility in managing compute resources. They can also be resized at any time, even while running, to accommodate varying workloads<sup>910</sup>. References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 243**

Which privilege allows a user to change the state of a virtual warehouse that was set to AUTO\_RESUME = FALSE?

- A. APPLYBUDGET
- B. USAGE
- C. MONITOR
- D. OPERATE

**Answer: D** (LEAVE A REPLY)

**NEW QUESTION: 244**

Which of the following features, associated with Continuous Data Protection (CDP), require additional Snowflake-provided data storage? (Choose two.)

- A. Tri-Secret Secure
- B. Time Travel
- C. Fail-safe
- D. Data encryption
- E. External stages

**Answer: B,C** (LEAVE A REPLY)

The features associated with Continuous Data Protection (CDP) that require additional Snowflake-provided data storage are Time Travel and Fail-safe. Time Travel allows users to access historical data within a defined period, while Fail-safe provides an additional layer of data protection beyond the Time Travel period. References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 245**

What action should be taken if a large number of concurrent queries are queued in a virtual warehouse?

- A. Enable auto-resume on the warehouse.
- B. Scale-up by resizing the warehouse.
- C. Scale-out with a multi-cluster warehouse.
- D. Disable auto-suspend on the warehouse.

**Answer: B** ([LEAVE A REPLY](#))

**NEW QUESTION: 246**

How can a dropped internal stage be restored?

- A. Enable Time Travel.
- B. Clone the dropped stage.
- C. Execute the UNDROP command.
- D. Recreate the dropped stage.

**Answer: D** ([LEAVE A REPLY](#))

Once an internal stage is dropped in Snowflake, it cannot be recovered or restored using Time Travel or UNDROP commands. The only option is to recreate the dropped stage

**NEW QUESTION: 247**

Which Snowflake privilege is required on a pipe object to pause or resume pipes?

- A. OPERATE
- B. READ
- C. SELECT
- D. USAGE

**Answer: (SHOW ANSWER)**

OPERATE. In Snowflake, to pause or resume a pipe, the OPERATE privilege is required on the pipe object. The OPERATE privilege allows users to perform operational tasks on specific objects such as pipes, tasks, and streams. Specifically, for a pipe, the OPERATE privilege enables the user to execute the ALTER PIPE ... SET PIPE\_EXECUTION\_PAUSED=TRUE or ALTER PIPE ... SET PIPE\_EXECUTION\_PAUSED=FALSE commands, which are used to pause or resume the pipe, respectively.

Here's a step-by-step explanation and reference:

Understanding Pipe Operations: Pipes in Snowflake are used for continuous data loading from staging areas into Snowflake tables. Managing pipes involves operations such as creating, monitoring, pausing, and resuming.

Privileges for Pipe Operations: The OPERATE privilege is essential for pausing and resuming pipes. This privilege is more specific than general object access privileges like SELECT or USAGE and is tailored for operational control.

Granting the OPERATE Privilege: To grant the OPERATE privilege on a pipe, an administrator or a user with the necessary grants can execute the SQL command:

```
GRANT OPERATE ON PIPE <pipe_name> TO ROLE <role_name>;
```

Pausing and Resuming Pipes: Once the OPERATE privilege is granted, the user or role can pause the pipe using:

```
ALTER PIPE <pipe_name> SET PIPE_EXECUTION_PAUSED=TRUE;
```

To resume the pipe, they use:

```
ALTER PIPE <pipe_name> SET PIPE_EXECUTION_PAUSED=FALSE;
```

### NEW QUESTION: 248

The following settings are configured:

THE MIN\_DATA\_RETENTION\_TIME\_IN\_DAYS is set to 5 at the account level. THE DATA\_RETENTION\_TIME\_IN\_DAYS is set to 2 at the object level.

For how many days will the data be retained at the object level?

- A. 2
- B. 3
- C. 5
- D. 7

**Answer: A** ([LEAVE A REPLY](#))

The settings shown in the image indicate that the data retention time in days is configured at two different levels: the account level and the object level. At the account level, the MIN\_DATA\_RETENTION\_TIME\_IN\_DAYS is set to 5 days, and at the object level, the DATA\_RETENTION\_TIME\_IN\_DAYS is set to 2 days. Since the object level setting has a lower value, it takes precedence over the account level setting for the specific object. Therefore, the data will be retained for 2 days at the object level. References: Snowflake Documentation on Data Retention Policies

### NEW QUESTION: 249

What is the impact of removing some of the privileges granted to a role?

- A. The role will retain all privileges until a new session is started.
- B. The role will lose access to objects that are not granted to the updated role.
- C. The user will have to reinstate the role using the USE ROLE command after the change.
- D. Any users who are actively using the role will have their current session terminated.

**Answer: (**[SHOW ANSWER](#)**)**

The correct answer is B. The role will lose access to objects that are not granted to the updated role .

Snowflake access control is based on privileges granted to roles. Users receive access to objects through the active role and any inherited roles in the role hierarchy. If privileges are revoked from a role, that role no longer has those privileges unless they are still inherited through another granted role.

Why B is correct:

When privileges are removed from a role, users using that role lose the ability to perform actions that depended on the revoked privileges. Access is controlled by the current set of grants assigned to the role.

Why the other options are incorrect:

- A). A role does not retain revoked privileges until a new session starts. Revoked privileges affect access according to Snowflake access control rules.
- C). The user does not need to reissue USE ROLE just because privileges were revoked.
- D). Snowflake does not terminate active sessions simply because a privilege is removed from a role.

Official Snowflake documentation reference:

Snowflake documentation explains that privileges are granted to roles, roles are granted to users or other roles, and revoking privileges removes the ability to perform the associated actions unless access exists through another role path.

Reference: Snowflake Documentation - Access control overview; Snowflake Documentation - GRANT privileges; Snowflake Documentation - REVOKE privileges; SnowPro Core Study Guide - Security and Access Control.

#### NEW QUESTION: 250

What is the expiration period for a file URL used to access unstructured data in cloud storage?

- A. The length of time specified in the `expiration_time` argument
- B. The same length of time as the expiration period for the query results cache
- C. An unlimited amount of time
- D. The remainder of the session

**Answer: B** ([LEAVE A REPLY](#))

#### NEW QUESTION: 251

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

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

**Answer: A,C** ([LEAVE A REPLY](#))

A clustering key can be defined when a table is created by appending a `CLUSTER` Where each clustering key consists of one or more table columns/expressions, which can be of any data type, except `GEOGRAPHY`, `VARIANT`, `OBJECT`, or `ARRAY`

<https://docs.snowflake.com/en/user-guide/tables-clustering-keys>

#### NEW QUESTION: 252

When is a DataFrame processed in Snowpark?

- A. When the DataFrame is created
- B. When a `collect()` action is triggered
- C. When there is sufficient memory available in the virtual warehouse
- D. When a new transformation is applied on the DataFrame

**Answer: B** ([LEAVE A REPLY](#))

The correct answer is B. When a `collect()` action is triggered .

Snowpark DataFrames are evaluated lazily. Creating a DataFrame or applying transformations does not immediately execute the query in Snowflake. Execution happens only when an action is called.

Why B is correct:

`collect()` is an action. When it is called, Snowpark sends the accumulated DataFrame logic to Snowflake for execution and returns the results to the client.

Example:

```
df = session.table( " CUSTOMERS " )
filtered_df = df.filter(df[ " REGION " ] == " WEST " )
# Query is processed when this action runs:
rows = filtered_df.collect()
```

Why the other options are incorrect:

- A). Creating a DataFrame builds a logical plan but does not process the data.
- C). Warehouse memory availability is not what triggers DataFrame execution.
- D). Applying a transformation adds to the logical plan, but it does not execute the plan.

Official Snowflake documentation reference:

Snowflake documentation describes Snowpark DataFrames as lazily evaluated. Transformations are not executed until an action, such as `collect()`, is called.

Reference: Snowflake Documentation - Snowpark DataFrames; Snowflake Documentation - Snowpark actions and transformations; SnowPro Core Study Guide - SQL and Snowflake Objects.

### NEW QUESTION: 253

A Snowflake Practitioner manages a database called SALES\_DB. This database has a PUBLIC schema and an ORDERS table that is updated frequently. The database has a data retention period of 7 days. On November 16, 2025, at 3:00 PM, the Practitioner accidentally ran a DELETE command that removed all records from the ORDERS table. Which statement will recover all deleted records?

- A. `CREATE TABLE orders_restored CLONE orders AT (TIMESTAMP = > ' 2025-11-16 10:00:00 ' ::TIMESTAMP);`
- B. `CREATE TABLE orders_restored CLONE orders BEFORE (TIMESTAMP = > ' 2025-11-16 15:00:00' ::TIMESTAMP);`
- C. `CREATE TABLE orders_restored CLONE orders BEFORE (OFFSET = > -60*5);`
- D. `CREATE TABLE orders_restored CLONE orders AT (TIMESTAMP = > ' 2025-11-16 15:00:00 ' ::TIMESTAMP);`

**Answer: B (LEAVE A REPLY)**

To recover data deleted at 3:00 PM, the clone should be created from a point immediately before the delete statement. `BEFORE (TIMESTAMP = > ' 2025-11-16 15:00:00 ' )` returns the table state before that timestamp.

Why B is correct:

Snowflake Time Travel supports cloning a table at or before a specific timestamp within the retention period.

Reference: Snowflake Documentation - Time Travel; `CREATE TABLE ... CLONE; AT | BEFORE`.

==

### NEW QUESTION: 254

Which command will consume credits from a virtual warehouse?

- A. `SELECT MAX(UUID) FROM SNOWFLAKE;`
- B. `SELECT COUNT(FLAKE_ID) FROM SNOWFLAKE GROUP BY UUID;`
- C. `SHOW TABLES LIKE ' SNOWFL%' IN SCHEMA( ' PUBLIC ' );`
- D. `SELECT COUNT(*) FROM SNOWFLAKE;`

**Answer: (SHOW ANSWER)**

The correct answers are A, B, and D , assuming SNOWFLAKE is a table or queryable object.

Important exam-note correction:

This question is written as a single-answer question, but more than one listed command can consume virtual warehouse credits. Any SELECT query that scans or processes table data generally requires a running virtual warehouse and can consume warehouse credits.

Why A, B, and D are correct:

These are SELECT statements against a table-like object. Query execution that scans or aggregates data uses compute resources from a virtual warehouse.

Examples:

```
SELECT MAX(UUID) FROM SNOWFLAKE;
```

```
SELECT COUNT(FLAKE_ID) FROM SNOWFLAKE GROUP BY UUID;
```

```
SELECT COUNT(*) FROM SNOWFLAKE;
```

These involve SQL query processing and therefore require warehouse compute unless the result is fully served from cache.

Why C is incorrect:

SHOW commands are metadata commands handled by Snowflake's cloud services layer. They do not require an active virtual warehouse in the same way table-scanning SELECT queries do.

Official Snowflake documentation reference:

Snowflake documentation explains that virtual warehouses provide the compute resources required to execute SQL queries and DML operations. Metadata commands such as SHOW are handled by the cloud services layer.

Reference: Snowflake Documentation - Virtual warehouses; Snowflake Documentation - SHOW commands; SnowPro Core Study Guide - Snowflake Architecture and Resource Management.

#### **NEW QUESTION: 255**

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** ([LEAVE A REPLY](#))

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 structures<sup>1</sup>.

#### **NEW QUESTION: 256**

Which ACCOUNT\_USAGE view can be used to identify the masking policy assigned to an object?

- A.** ACCESS\_HISTORY
- B.** OBJECT\_DEPENDENCIES
- C.** TAG\_REFERENCES
- D.** POLICY\_REFERENCES

**Answer: D** ([LEAVE A REPLY](#))

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been**

**corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 257

In Snowpipe, what is the default behavior of the ON\_ERROR copy option?

- A. CONTINUE
- B. SKIP\_FILE
- C. SKIP\_FILE\_ < num >
- D. ABORT\_STATEMENT

**Answer: B (LEAVE A REPLY)**

The correct answer is B. SKIP\_FILE .

Important correction:

The provided answer A. CONTINUE is not correct for Snowpipe.

For regular bulk loading with COPY INTO < table > , the default ON\_ERROR behavior is commonly associated with ABORT\_STATEMENT. However, Snowpipe has different behavior. In Snowpipe, the default ON\_ERROR option is SKIP\_FILE .

Why B is correct:

With ON\_ERROR = SKIP\_FILE, if Snowpipe encounters an error while loading a file, the file is skipped instead of causing the entire continuous loading process to fail.

Why the other options are incorrect:

- A). CONTINUE continues loading and skips individual error rows, but it is not the default Snowpipe behavior.
- C). SKIP\_FILE\_ < num > skips a file after a specified number of errors, but it is not the default.
- D). ABORT\_STATEMENT is associated with standard COPY INTO < table > behavior, not the default Snowpipe behavior.

Official Snowflake documentation reference:

Snowflake documentation for Snowpipe explains that the default ON\_ERROR copy option for Snowpipe is SKIP\_FILE.

Reference: Snowflake Documentation - Snowpipe; Snowflake Documentation - CREATE PIPE; Snowflake Documentation - COPY INTO < table > copy options; SnowPro Core Study Guide - Data Loading and Unloading.

#### NEW QUESTION: 258

Which column data type is supported when defining a clustering key on a table?

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

**Answer: A (LEAVE A REPLY)**

The correct answer is A. VARCHAR .

A clustering key can be defined on supported table columns or expressions to help Snowflake organize micro- partitions and improve pruning for certain query patterns. Standard scalar data types such as VARCHAR are supported.

Why A is correct:

VARCHAR is a standard scalar data type and can be used in a clustering key.

Example:

CREATE TABLE customer\_orders (

```
customer_id VARCHAR,  
order_date DATE,  
amount NUMBER  
)
```

```
CLUSTER BY (customer_id);
```

Why the other options are incorrect:

B). VARIANT stores semi-structured data. Clustering directly on raw semi-structured columns is not the best answer here; clustering keys should use supported scalar columns or expressions.

C). OBJECT is a semi-structured data type and is not the correct choice for a clustering key column in this question.

D). GEOGRAPHY is not the correct supported clustering key data type in this context.

Official Snowflake documentation reference:

Snowflake documentation explains that clustering keys can be defined using table columns or expressions, and standard scalar columns such as string, number, and date/time columns are commonly used for clustering.

Reference: Snowflake Documentation - Clustering keys and clustered tables; Snowflake Documentation - Table data types; SnowPro Core Study Guide - Performance Optimization.

### NEW QUESTION: 259

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

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

**Answer: C,D ([LEAVE A REPLY](#))**

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

### NEW QUESTION: 260

Which command should be used to load data incrementally based on column values that are specified in the source table or subquery?

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

**Answer: A ([LEAVE A REPLY](#))**

The MERGE command in Snowflake is used for incremental loading based on column values in a source table or subquery. It enables the insertion, updating, or deletion of records in a target table depending on whether matching rows are found, making it ideal for loading data that changes incrementally, such as daily updates or modifications.

### NEW QUESTION: 261

A table, STUDENTS, is created using this command:

An example value for the marks is:

Which query will return all of the students ' names and history marks?

- A. SELECT NAME, MARKS [0] FROM STUDENTS;
- B. SELECT NAME, MARKS: HISTORY FROM STUDENTS;
- C. SELECT NAME, MARKS:history FROM STUDENTS;
- D. SELECT NAME, GET (MARKS, ' history ' ) FROM STUDENTS;

**Answer:** ([SHOW ANSWER](#))

#### NEW QUESTION: 262

When publishing a Snowflake Data Marketplace listing into a remote region what should be taken into consideration? (Choose two.)

- A. There is no need to have a Snowflake account in the target region, a share will be created for each user.
- B. The listing is replicated into all selected regions automatically, the data is not.
- C. The user must have the ORGADMIN role available in at least one account to link accounts for replication.
- D. Shares attached to listings in remote regions can be viewed from any account in an organization.
- E. For a standard listing the user can wait until the first customer requests the data before replicating it to the target region.

**Answer:** B,C ([LEAVE A REPLY](#))

When publishing a Snowflake Data Marketplace listing into a remote region, it's important to note that while the listing is replicated into all selected regions automatically, the data itself is not. Therefore, the data must be replicated separately. Additionally, the user must have the ORGADMIN role in at least one account to manage the replication of accounts<sup>1</sup>.

#### NEW QUESTION: 263

What are valid values for the FIELD\_OPTIONALLY\_ENCLOSED\_BY option in the copy into < location > command used during data unloading? (Select TWO).

- A. ' NULL '
- B. ' NONE '
- C. NULL
- D. NONE
- E. Single quote character ( ' )

**Answer:** ([SHOW ANSWER](#))

#### NEW QUESTION: 264

What are potential impacts of storing non-native values like dates and timestamps in a variant column in Snowflake?

- A. Faster query performance and increased storage consumption
- B. Slower query performance and increased storage consumption
- C. Faster query performance and decreased storage consumption
- D. Slower query performance and decreased storage consumption

**Answer:** B ([LEAVE A REPLY](#))

Storing non-native values, such as dates and timestamps, in a VARIANT column in Snowflake can lead to slower query performance and increased storage consumption. VARIANT is a semi-structured data type that allows storing JSON, AVRO, ORC, Parquet, or XML data in a single column. When non-native data types are stored as VARIANT, Snowflake must perform implicit conversion to process these values, which can slow down query execution. Additionally, because the VARIANT data type is designed to accommodate a wide variety of data formats, it often requires more storage space compared to storing data in native, strongly-typed columns that are optimized for specific data types.

The performance impact arises from the need to parse and interpret the semi-structured data on the fly during query execution, as opposed to directly accessing and operating on optimally stored data in its native format. Furthermore, the increased storage consumption is a result of the overhead associated with storing data in a format that is less space-efficient than the native formats optimized for specific types of data.

References:

Snowflake Documentation on Semi-Structured Data: Semi-Structured Data

### NEW QUESTION: 265

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 (LEAVE A REPLY)**

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-C03 ] SnowPro Core Certification Exam Study Guide

### NEW QUESTION: 266

A table gets updated throughout the day. A user wants to automate the process of table changes by capturing only the changed rows every 20 minutes and applying transformations without manual intervention. How can this be accomplished?

- A. Create a stream and run a task only when there are changes.
- B. Create a task that triggers every 20 minutes.
- C. Create a stream and schedule a task every 20 minutes.
- D. Use a dynamic table to materialize changes and run a task for the business transformations every 20 minutes.

**Answer: C (LEAVE A REPLY)**

The correct answer is C. Create a stream and schedule a task every 20 minutes.

A stream is used to capture change data capture, or CDC, information for a table. It records inserts, updates, and deletes made to the source table. A task is used to automate SQL execution on a schedule. Together, streams and tasks are commonly used to automate incremental data processing in Snowflake.

Why C is correct:

The requirement has two parts:

Capture only changed rows.

Run the transformation automatically every 20 minutes.

A stream satisfies the first requirement because it tracks changed rows. A scheduled task satisfies the second requirement because it can run SQL logic at a defined interval, such as every 20 minutes.

Example concept:

```
CREATE STREAM my_stream
ON TABLE source_table;
CREATE TASK my_task
WAREHOUSE = my_warehouse
SCHEDULE = ' 20 MINUTE '
AS
INSERT INTO target_table
SELECT ...
FROM my_stream;
```

Why the other options are incorrect:

- A). This is partially correct, but it does not clearly satisfy the requirement to run every 20 minutes. A task can include a `WHEN SYSTEM$STREAM_HAS_DATA(...)` condition, but the question specifically asks for processing every 20 minutes.
- B). A task alone can run every 20 minutes, but it does not capture only changed rows.
- D). Dynamic tables are used for declarative automated data transformation, but the question specifically asks to capture only changed rows from a frequently updated table, which is the purpose of a stream.

Official Snowflake documentation reference:

Snowflake documentation describes streams as objects that record data manipulation language, or DML, changes made to tables. Snowflake tasks execute SQL statements on a schedule or as part of a task graph.

Reference: Snowflake Documentation - Streams; Snowflake Documentation - Introduction to tasks; SnowPro Core Study Guide - SQL and Snowflake Objects.

### NEW QUESTION: 267

Which function returns the URL of a stage using the stage name as the input?

- A. `BUILD_STAGE_FILE_URL`
- B. `BUILD_SCOPED_FILE_URL`
- C. `GET_PREIGNED_URL`
- D. `GET STAGE LOCATION`

**Answer: C ([LEAVE A REPLY](#))**

The function in Snowflake that returns the URL of a stage using the stage name as the input is C. `GET_PREIGNED_URL`.

This function generates a pre-signed URL for a specific file in a stage, enabling secure, temporary access to that file without requiring Snowflake credentials. While the function is primarily used for accessing files in external stages, such as Amazon S3 buckets, it is instrumental in scenarios requiring direct, secure file access for a limited time.

It's important to note that as of my last update, Snowflake's documentation does not specifically list a function named `GET_PREIGNED_URL` for directly obtaining a stage's URL by its name. The description aligns closely with functionality available in cloud storage services (e.g., AWS S3's presigned URLs) which can be used in conjunction with Snowflake stages for secure, temporary access to files. For direct interaction with stages and their files, Snowflake offers various functions and commands, but the exact match for generating a presigned URL through a simple function call may vary or require leveraging external cloud services APIs in addition to Snowflake's capabilities.

References:

Snowflake Documentation and cloud services (AWS, Azure, GCP) documentation on presigned URLs and stage interactions.

#### **NEW QUESTION: 268**

What function can be used with the recursive argument to return a list of distinct key names in all nested elements in an object?

- A.** FLATTEN
- B.** GET\_PATH
- C.** CHECK\_JSON
- D.** PARSE\_JSON

**Answer: A** ([LEAVE A REPLY](#))

The FLATTEN function can be used with the recursive argument to return a list of distinct key names in all nested elements within an object. This function is particularly useful for working with semi-structured data in Snowflake

#### **NEW QUESTION: 269**

Which Snowflake technique can be used to improve the performance of a query?

- A.** Clustering
- B.** Indexing
- C.** Fragmenting
- D.** Using INDEX\_\_HINTS

**Answer: A** ([LEAVE A REPLY](#))

Clustering is a technique used in Snowflake to improve the performance of queries. It involves organizing the data in a table into micro-partitions based on the values of one or more columns. This organization allows Snowflake to efficiently prune non-relevant micro-partitions during a query, which reduces the amount of data scanned and improves query performance.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Clustering

#### **NEW QUESTION: 270**

What feature specifies which roles are allowed to see a particular record in a table?

- A.** Dynamic Data Masking
- B.** Secure Data Sharing
- C.** Row access policy
- D.** Object tagging

**Answer: C** ([LEAVE A REPLY](#))

The correct answer is C. Row access policy .

A row access policy is used to implement row-level security in Snowflake. It controls whether a row is visible in query results based on conditions such as the current role, current user, or values in a mapping table.

Why C is correct:

The question asks which feature controls visibility of particular records, meaning rows. Row access policies are designed exactly for that purpose.

Example concept:

```
CREATE ROW ACCESS POLICY region_policy
AS (region STRING) RETURNS BOOLEAN - >
region IN (
SELECT allowed_region
FROM role_region_mapping
WHERE role_name = CURRENT_ROLE()
);
```

This policy can allow users with different roles to see only the rows assigned to their region.

Why the other options are incorrect:

- A). Dynamic Data Masking controls how column values are displayed. It masks data in columns but does not filter entire rows.
- B). Secure Data Sharing allows data to be shared between accounts, but it is not the row-level filtering feature itself.
- D). Object tagging classifies or labels objects for governance, but it does not directly control row visibility.

Official Snowflake documentation reference:

Snowflake documentation describes row access policies as schema-level objects that determine whether a row is visible in query results.

Reference: Snowflake Documentation - Row access policies; Snowflake Documentation - Row-level security; SnowPro Core Study Guide - Data Protection and Governance.

==

### NEW QUESTION: 271

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: (SHOW ANSWER)**

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

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

### NEW QUESTION: 272

What information does the Query Profile provide?

- A. Graphical representation of the data model
- B. Statistics for each component of the processing plan
- C. Detailed Information about the database schema
- D. Real-time monitoring of the database operations

**Answer: B** ([LEAVE A REPLY](#))

The Query Profile in Snowflake provides a graphical representation and statistics for each component of the query's execution plan. This includes details such as the execution time, the number of rows processed, and the amount of data scanned for each operation within the query. The Query Profile is a crucial tool for understanding and optimizing the performance of queries, as it helps identify potential bottlenecks and inefficiencies.

References:

Snowflake Documentation: Understanding the Query Profile

### NEW QUESTION: 273

What is the primary purpose of a directory table in Snowflake?

- A. To store actual data from external stages
- B. To automatically expire file URLs for security
- C. To manage user privileges and access control
- D. To store file-level metadata about data files in a stage

**Answer: D** ([LEAVE A REPLY](#))

A directory table in Snowflake is used to store file-level metadata about the data files in a stage. It is conceptually similar to an external table and provides information such as file size, last modified timestamp, and file URL. References: [COF-C02]

SnowPro Core Certification Exam Study Guide

### NEW QUESTION: 274

Which property or parameter can be used to temporarily disable Multi-Factor Authentication (MFA) for a Snowflake user?

- A. ALLOW\_CLIENT\_MFA\_CACHING
- B. EXT\_AUTHN\_DUO
- C. DISABLE\_MFA
- D. MINS\_TO\_BYPASS\_MFA

**Answer:** ([SHOW ANSWER](#))

### NEW QUESTION: 275

Which chart type is supported in Snowsight for Snowflake users to visualize data with dashboards?

- A. Area chart
- B. Box plot
- C. Heat grid
- D. Pie chart

**Answer: A** ([LEAVE A REPLY](#))

Snowsight, Snowflake's user interface for exploring, analyzing, and visualizing data, supports a variety of chart types for creating dashboards and visualizations. One of the supported chart types in Snowsight is the Area Chart (A). Area charts are

useful for representing quantities over time and can be used to highlight volume change and rate of change, as well as to compare multiple quantities.

While Snowsight supports many types of visualizations to help users analyze their data effectively, including line charts, bar charts, and scatter plots, it ' s important to select the specific reference documentation or release notes for the most current list of supported chart types, as Snowflake continues to enhance and update Snowsight ' s capabilities.

As of the last update, Box plots (B), Heat grids (C), and Pie charts (D) are types of visualizations that may be supported in various analytics and visualization tools, but for the specific context of Snowsight ' s currently confirmed features, Area charts are a verified option for users to visualize their data.

Reference: Snowflake Documentation on Snowsight (<https://docs.snowflake.com/en/user-guide/ui-snowsight.html>)

### NEW QUESTION: 276

What can be shared when the `SECURE_OBJECTS_ONLY` property is set to `FALSE`?

- A. View
- B. Stored procedure
- C. Storage integration access
- D. User-Defined Function, or UDF

**Answer: A ([LEAVE A REPLY](#))**

The correct answer is A. View .

In Snowflake data sharing, secure objects are normally required when sharing views or functions because secure objects prevent sensitive underlying query details from being exposed. However, when the `SECURE_OBJECTS_ONLY` property is set to `FALSE`, non-secure views can be shared.

Why A is correct:

A view can be shared when `SECURE_OBJECTS_ONLY = FALSE`. This setting relaxes the requirement that only secure views be included.

Why the other options are incorrect:

- B). Stored procedures are not standard objects shared through Snowflake secure data sharing in the same way as tables and views.
- C). Storage integration access is not shared as a database object through a Snowflake share.
- D). UDF sharing has security restrictions, and the safest exam answer here is a view when `SECURE_OBJECTS_ONLY` is disabled.

Official Snowflake documentation reference:

Snowflake documentation explains that shares can include database objects such as tables and views. Secure object requirements apply to certain shared objects, and the `SECURE_OBJECTS_ONLY` property controls whether only secure objects are allowed.

Reference: Snowflake Documentation - Secure Data Sharing; Snowflake Documentation - Shares; Snowflake Documentation - `SECURE_OBJECTS_ONLY`; SnowPro Core Study Guide - Data Protection and Governance.

### NEW QUESTION: 277

Which copy INTO command outputs the data into one file?

- A. `SINGLE=TRUE`

- B. MAX\_FILE\_NUMBER=1
- C. FILE\_NUMBER=1
- D. MULTIPLE=FAISE

**Answer: B (LEAVE A REPLY)**

The COPY INTO command in Snowflake can be configured to output data into a single file by setting the MAX\_FILE\_NUMBER option to 1. This option limits the number of files generated by the command, ensuring that only one file is created regardless of the amount of data being exported.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Data Unloading

### NEW QUESTION: 278

A user is setting up Snowpipe automation to load data from an Amazon S3 bucket into Snowflake. The user wants Snowflake to automatically start loading files as soon as new ones are added to the bucket. Which steps are required to configure this automation using cloud messaging? Select TWO.

- A. Configure an S3 event notification to an Amazon Simple Notification Service, or SNS, topic.
- B. Create a task that monitors the stage for new files.
- C. Call the insertFiles REST API endpoint after each file upload.
- D. Enable AUTO\_INGEST on the Snowpipe definition.
- E. Define an external table and use it as the source for a materialized view.

**Answer: A,D (LEAVE A REPLY)**

For Snowpipe auto-ingest on Amazon S3, cloud event notifications are configured so Snowflake is notified when new files arrive. The pipe must also be created with AUTO\_INGEST = TRUE.

Why A and D are correct:

Amazon S3 event notifications can publish file arrival events through Amazon SNS/SQS integration, and Snowpipe uses those notifications when AUTO\_INGEST is enabled.

Reference: Snowflake Documentation - Snowpipe auto-ingest on Amazon S3; CREATE PIPE.

==

### NEW QUESTION: 279

What are streams used for in Snowflake?

- A. To execute data pipelines
- B. To track DML changes that have been made to tables
- C. To define the data file formats being loaded into Snowflake
- D. To cache data to be used to improve future query performance

**Answer: B (LEAVE A REPLY)**

The correct answer is B. To track DML changes that have been made to tables .

A stream is a Snowflake object used for change data capture. It records changes made to a source table so those changes can be consumed later by SQL logic, usually as part of an incremental processing pipeline.

Why B is correct:

Streams track DML changes such as:

INSERT  
UPDATE  
DELETE

They also expose metadata columns such as:

METADATA\$ACTION  
METADATA\$ISUPDATE  
METADATA\$ROW\_ID

These columns help identify the type of change and the affected row.

Why the other options are incorrect:

- A). Tasks execute scheduled SQL and can be used to run data pipelines. Streams track the changes consumed by those pipelines.
- C). File formats define how staged files are interpreted during loading or unloading. Streams are unrelated to file format definitions.
- D). Caching is handled by Snowflake result cache and warehouse cache mechanisms, not by streams.

Official Snowflake documentation reference:

Snowflake documentation describes streams as objects that record DML changes made to tables, including inserts, updates, and deletes.

Reference: Snowflake Documentation - Streams; Snowflake Documentation - Change data capture; SnowPro Core Study Guide - SQL and Snowflake Objects.

==

### NEW QUESTION: 280

What information is included in the output of an EXPLAIN\_PLAN?

- A. The predicted time it will take to run the operation.
- B. The number of credits that the operation will consume.
- C. The type of data that will be returned when the operation is run.
- D. The total number of micro-partitions that the operation will reference.

**Answer: D (LEAVE A REPLY)**

The correct answer is D. The total number of micro-partitions that the operation will reference .

In Snowflake, EXPLAIN or EXPLAIN USING returns the logical execution plan for a SQL statement. The plan can include information about the operations Snowflake expects to perform, such as table scans, filters, joins, and estimated partitions involved in the query.

Why D is correct:

The explain plan can show partition-related information for table scan operations, including the number of partitions that may be scanned. This is useful for understanding how well Snowflake can prune micro- partitions.

Why the other options are incorrect:

- A). EXPLAIN\_PLAN does not provide a guaranteed predicted query runtime. Actual execution time depends on warehouse size, load, caching, pruning, and other runtime factors.
- B). EXPLAIN\_PLAN does not directly state the number of credits the query will consume. Credit consumption is based mainly on warehouse runtime and size.
- C). EXPLAIN\_PLAN describes the execution plan, not the actual result data type returned by the query.

Official Snowflake documentation reference:

Snowflake documentation explains that the EXPLAIN command displays the logical execution plan for a SQL statement. Query plans can include details about operations such as scans, joins, filters, and partitions.

Reference: Snowflake Documentation - EXPLAIN; Snowflake Documentation - Query Profile and query optimization; SnowPro Core Study Guide - Performance Optimization.

==

#### **NEW QUESTION: 281**

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

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

**Answer: C** ([LEAVE A REPLY](#))

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

#### **NEW QUESTION: 282**

Which use case does the search optimization service support?

- A.** Disjuncts (OR) in join predicates
- B.** LIKE/ILIKE/RLIKE join predicates
- C.** Join predicates on VARIANT columns
- D.** Conjunctions (AND) of multiple equality predicates

**Answer: D** ([LEAVE A REPLY](#))

The search optimization service in Snowflake supports use cases involving conjunctions (AND) of multiple equality predicates. This service enhances the performance of queries that include multiple equality conditions by utilizing search indexes to quickly filter data without scanning entire tables or partitions. It is particularly beneficial for improving the response times of complex queries that rely on specific data matching across multiple conditions.

References:

Snowflake Documentation: Search Optimization Service

#### **NEW QUESTION: 283**

What are the possible values within a METADATAACTION column in a Snowflake stream? (Select TWO).

- A.** INSERT
- B.** UPDATE
- C.** DELETE
- D.** TRUNCATE
- E.** UPSERT

**Answer: A,C** ([LEAVE A REPLY](#))

In Snowflake streams, the METADATASACTION column indicates the type of data manipulation operation that has occurred. The possible values include INSERT and DELETE.

INSERT: Indicates that a new row has been inserted into the table.

DELETE: Indicates that a row has been deleted from the table.

References:

Snowflake Documentation: Change Data Capture (CDC) with Streams

Snowflake Documentation: Stream Data Capture

#### **NEW QUESTION: 284**

Which type of charts are supported by Snowsight? {Select TWO}

- A. Gantt charts
- B. Scatterplots
- C. Pie charts
- D. Line charts
- E. Flowcharts

**Answer:** ([SHOW ANSWER](#))

#### **NEW QUESTION: 285**

When unloading to a stage, which of the following is a recommended practice or approach?

- A. Set SINGLE: = true for larger files
- B. Use OBJECT\_CONSTRUCT ( \* ) when using Parquet
- C. Avoid the use of the CAST function
- D. Define an individual file format

**Answer: D** ([LEAVE A REPLY](#))

When unloading data to a stage, it is recommended to define an individual file format. This ensures that the data is unloaded in a consistent and expected format, which can be crucial for downstream processing and analysis.

#### **NEW QUESTION: 286**

When resizing a currently running virtual warehouse, what happens to in-progress queries?

- A. In-progress queries will be cancelled.
- B. Queries will continue running, using the new warehouse size.
- C. Queries will continue running, using the current warehouse size.
- D. Queries will be paused and automatically reprocessed using the new warehouse size.

**Answer: C** ([LEAVE A REPLY](#))

The correct answer is C. Queries will continue running, using the current warehouse size .

When a virtual warehouse is resized while it is running, the resize does not affect queries that are already in progress. Running queries continue using the warehouse resources that were available when they started. The new warehouse size applies to queries that start after the resize operation.

Why C is correct:

Snowflake does not cancel or restart existing queries when a warehouse is resized. In-progress queries continue executing with the current warehouse size.

Why the other options are incorrect:

- A). Resizing a warehouse does not automatically cancel in-progress queries.
- B). In-progress queries do not switch to the new warehouse size.
- D). Queries are not paused and reprocessed after resizing.

Official Snowflake documentation reference:

Snowflake documentation explains that resizing a running warehouse affects only new queries and queued queries. Currently running queries continue to use the warehouse size that was in effect when they started.

Reference: Snowflake Documentation - Resizing a warehouse; Snowflake Documentation - Virtual warehouses; SnowPro Core Study Guide - Snowflake Account and Resource Management.

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 287

Queries are frequently run on very large tables and return only a few rows.

Use of which feature or service will optimize the query performance?

- A. Automatic Clustering
- B. The Search Optimization Service
- C. The Query Acceleration Service
- D. A Materialized View

**Answer: B (LEAVE A REPLY)**

Search Optimization Service enhances performance for selective queries that return a small number of rows from large datasets. It avoids full table scans.

Reference:

Snowflake Docs: Search Optimization Service

#### NEW QUESTION: 288

Which virtual warehouse consideration can help lower compute resource credit consumption?

- A. Setting up a multi-cluster virtual warehouse
- B. Resizing the virtual warehouse to a larger size
- C. Automating the virtual warehouse suspension and resumption settings
- D. Increasing the maximum cluster count parameter for a multi-cluster virtual warehouse

**Answer: C (LEAVE A REPLY)**

One key strategy to lower compute resource credit consumption in Snowflake is by automating the suspension and resumption of virtual warehouses. Virtual warehouses consume credits when they are running, and managing their operational times effectively can lead to significant cost savings.

A . Setting up a multi-cluster virtual warehouse increases parallelism and throughput but does not directly lower credit consumption. It is more about performance scaling than cost efficiency.

B . Resizing the virtual warehouse to a larger size increases the compute resources available for processing queries, which increases the credit consumption rate. This option does not help in lowering costs.

C . Automating the virtual warehouse suspension and resumption settings: This is a direct method to manage credit consumption efficiently. By automatically suspending a warehouse when it is not in use and resuming it when needed, you can avoid consuming credits during periods of inactivity. Snowflake allows warehouses to be configured to automatically suspend after a specified period of inactivity and to automatically resume when a query is submitted that requires the warehouse.

D . Increasing the maximum cluster count parameter for a multi-cluster virtual warehouse would potentially increase credit consumption by allowing more clusters to run simultaneously. It is used to scale up resources for performance, not to reduce costs.

Automating the operational times of virtual warehouses ensures that you only consume compute credits when the warehouse is actively being used for queries, thereby optimizing your Snowflake credit usage.

#### **NEW QUESTION: 289**

If a table was dropped and the `DATA_RETENTION_TIME_IN_DAYS` was set to 0, how can the data from the table be recovered?

- A. By creating a failover group
- B. By creating a Snowflake Support ticket
- C. By running the `UNDROP TABLE` command
- D. By running the `CLONE ... BEFORE` command

**Answer: B ([LEAVE A REPLY](#))**

The correct answer is B. By creating a Snowflake Support ticket .

If `DATA_RETENTION_TIME_IN_DAYS` = 0, Time Travel retention is disabled for the table. This means the table cannot be recovered using normal Time Travel operations such as `UNDROP TABLE` or `CLONE ... BEFORE`.

Why B is correct:

After Time Travel retention is unavailable, recovery may only be possible through Snowflake Fail-safe for permanent tables. Fail-safe is not directly accessible by users. Recovery from Fail-safe requires contacting Snowflake Support.

Important qualification:

This assumes the dropped table is a permanent table and the data is still within Fail-safe. Temporary and transient tables do not have Fail-safe, so they cannot be recovered this way.

Why the other options are incorrect:

- A). A failover group is used for replication and disaster recovery across accounts or regions. It does not recover a dropped table with no Time Travel retention.
- C). `UNDROP TABLE` relies on Time Travel retention. With retention set to 0, this is not available.
- D). `CLONE ... BEFORE` also relies on Time Travel. It cannot recover data when no retention period exists.

Official Snowflake documentation reference:

Snowflake documentation explains that Time Travel enables users to restore dropped objects within the retention period. Fail-safe is a separate recovery mechanism for permanent tables and is only available through Snowflake Support.

Reference: Snowflake Documentation - Time Travel; Fail-safe; UNDROP TABLE; SnowPro Core Study Guide - Data Protection and Governance.

==

### NEW QUESTION: 290

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** ([LEAVE A REPLY](#))

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: 291

Which constraint is enforced for standard Snowflake tables?

- A. NOT NULL
- B. PRIMARY KEY
- C. FOREIGN KEY
- D. UNIQUE

**Answer: (SHOW ANSWER)**

The correct answer is A. NOT NULL .

For standard Snowflake tables, Snowflake enforces the NOT NULL constraint. Other constraints such as PRIMARY KEY, FOREIGN KEY, and UNIQUE can be defined, but they are generally informational and are not enforced for standard tables.

Why A is correct:

A NOT NULL constraint prevents NULL values from being inserted into the constrained column.

Example:

```
CREATE TABLE customers (  
  customer_id NUMBER NOT NULL,  
  customer_name STRING  
);
```

Why the other options are incorrect:

- B). PRIMARY KEY constraints are not enforced for standard Snowflake tables.
- C). FOREIGN KEY constraints are not enforced for standard Snowflake tables.
- D). UNIQUE constraints are not enforced for standard Snowflake tables.

Official Snowflake documentation reference:

Snowflake documentation explains that for standard tables, NOT NULL constraints are enforced, while unique, primary key, and foreign key constraints are not enforced.

Reference: Snowflake Documentation - Constraints; Snowflake Documentation - Table constraints; SnowPro Core Study Guide - SQL and Snowflake Objects.

==

#### NEW QUESTION: 292

Which Snowflow object does not consume and storage costs?

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

**Answer:** ([SHOW ANSWER](#))

Temporary tables in Snowflake do not consume storage costs. They are designed for transient data that is needed only for the duration of a session. Data stored in temporary tables is held in the virtual warehouse's cache and does not persist beyond the session's lifetime, thereby not incurring any storage charges.

References:

Snowflake Documentation: Temporary Tables

#### NEW QUESTION: 293

Which command should be used to assign a key to a Snowflake user who needs to connect using key pair authentication?

- A. ALTER USER jsmith SET ENCRYPTED\_KEY='MIIBIjANBgkqh...';
- B. ALTER USER jsmith SET RSA\_P8\_KEY='MIIBIjANBgkqh...';
- C. ALTER USER jsmith SET RSA\_PRIVATE\_KEY='MIIBIjANBgkqh...';
- D. ALTER USER jsmith SET RSA\_PUBLIC\_KEY-MIIBIjANBgkqh...';

**Answer:** D ([LEAVE A REPLY](#))

#### NEW QUESTION: 294

In which hierarchy is tag inheritance possible?

- A. Organization → Account → Role
- B. Account → User → Schema
- C. Database → View → Column
- D. Schema → Table → Column

**Answer:** ([SHOW ANSWER](#))

In Snowflake, tag inheritance is a feature that allows tags, which are key-value pairs assigned to objects for the purpose of data governance and metadata management, to be inherited within a hierarchy. The hierarchy in which tag inheritance is possible is from Schema to Table to Column. This means that a tag applied to a schema can be inherited by the tables within that schema, and a tag applied to a table can be inherited by the columns within that table.

References: Snowflake Documentation on Tagging and Object Hierarchy

#### NEW QUESTION: 295

When using the COPY INTO < location > command, which parameter is used to unload table rows into separate files?

- A. DETAILED\_OUTPUT

- B. PARTITION BY**
- C. FILE\_FORMAT**
- D. HEADER**

**Answer: B (LEAVE A REPLY)**

The correct answer is B. PARTITION BY .

The PARTITION BY option in COPY INTO < location > partitions unloaded data into separate files based on the expression specified. This is commonly used to organize unloaded files by values such as date, region, or another column.

Why B is correct:

PARTITION BY controls how unloaded rows are grouped into different output paths/files.

Example:

```
COPY INTO @my_stage/sales/
```

```
FROM sales
```

```
PARTITION BY (TO_DATE(sale_date))
```

```
FILE_FORMAT = (TYPE = CSV);
```

This unloads rows into separate output paths based on the sale\_date expression.

Why the other options are incorrect:

A). DETAILED\_OUTPUT controls the amount of unload result information returned by the command.

C). FILE\_FORMAT controls the format of the unloaded files, such as CSV, JSON, or Parquet.

D). HEADER controls whether column headings are included in unloaded output files.

Official Snowflake documentation reference:

Snowflake documentation for COPY INTO < location > describes PARTITION BY as the option used to partition unloaded rows into separate files based on an expression.

Reference: Snowflake Documentation - COPY INTO < location > ; Snowflake Documentation - Unloading data with partitioning; SnowPro Core Study Guide - Data Loading and Unloading.

## **NEW QUESTION: 296**

What is the purpose of a stream in Snowflake?

- A.** To process and move data from multiple sources into Snowflake in real time
- B.** To provide a mechanism to deliver live or recorded media content to global cross-cloud users in real time
- C.** To record changes made to table data and metadata, so that the changed data can be processed
- D.** To process high volumes of data in a continuous, incremental manner using low-latency processing

**Answer: C (LEAVE A REPLY)**

The correct answer is C. To record changes made to table data and metadata, so that the changed data can be processed .

A Snowflake stream is a change tracking object. It records data manipulation language, or DML, changes made to a source table, including inserts, updates, and deletes. Streams are commonly used with tasks to build incremental data pipelines.

Why C is correct:

Streams allow users to identify what changed in a table since the stream was last consumed. This makes it possible to process only changed rows rather than reprocessing the entire table.

A stream exposes metadata columns such as:

```
METADATA$ACTION
```

```
METADATA$ISUPDATE
```

METADATA\$ROW\_ID

These metadata columns help identify the type of change and the affected row.

Why the other options are incorrect:

- A). This describes real-time ingestion, which is closer to Snowpipe, connectors, or streaming ingestion rather than streams.
- B). Snowflake streams are not media streaming services.
- D). This describes a general streaming data processing system, but a Snowflake stream itself does not process data. It records changes so other SQL processes can consume them.

Official Snowflake documentation reference:

Snowflake documentation describes streams as objects that record DML changes made to tables, including metadata about each change, so the changes can be consumed in downstream processing.

Reference: Snowflake Documentation - Streams; Snowflake Documentation - Stream metadata columns; SnowPro Core Study Guide - SQL and Snowflake Objects.

### NEW QUESTION: 297

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

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

**Answer: A ([LEAVE A REPLY](#))**

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

References:

Snowflake Documentation on Unloading Data

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

### NEW QUESTION: 298

How should the SPLIT\_TO\_TABLE( < string > , < delimiter > ) function be called?

- A. SELECT SPLIT\_TO\_TABLE < COL1, ' . ' > FROM DUAL;
- B. SELECT SPLIT\_TO\_TABLE( ' a.b.c ' , ' . ' );
- C. SELECT \* FROM TABLE(SPLIT\_TO\_TABLE( ' a.b.c ' , ' . ' ));
- D. SELECT \* FROM SPLIT\_TO\_TABLE( ' a-b.c ' , ' . ' );

**Answer: C ([LEAVE A REPLY](#))**

SPLIT\_TO\_TABLE() is a table function that splits a string into multiple rows based on a delimiter. To use it correctly in SQL, you must wrap it in a TABLE() clause.

Correct Syntax:

```
SELECT * FROM TABLE(SPLIT_TO_TABLE( ' a.b.c ' , ' . ' ));
```

Reference:

Snowflake Docs: SPLIT\_TO\_TABLE

### NEW QUESTION: 299

Which function returns the URL of a stage using the stage name as the input?

- A. BUILD\_STAGE\_FILE\_URL
- B. BUILD\_SCOPED\_FILE\_URL
- C. GET\_PREIGNED\_URL
- D. GET STAGE LOCATION

**Answer:** [\(SHOW ANSWER\)](#)

The function in Snowflake that returns the URL of a stage using the stage name as the input is C.

GET\_PREIGNED\_URL. This function generates a pre-signed URL for a specific file in a stage, enabling secure, temporary access to that file without requiring Snowflake credentials. While the function is primarily used for accessing files in external stages, such as Amazon S3 buckets, it is instrumental in scenarios requiring direct, secure file access for a limited time.

It's important to note that as of my last update, Snowflake's documentation does not specifically list a function named GET\_PREIGNED\_URL for directly obtaining a stage's URL by its name. The description aligns closely with functionality available in cloud storage services (e.g., AWS S3's presigned URLs) which can be used in conjunction with Snowflake stages for secure, temporary access to files. For direct interaction with stages and their files, Snowflake offers various functions and commands, but the exact match for generating a presigned URL through a simple function call may vary or require leveraging external cloud services APIs in addition to Snowflake's capabilities.

References:

Snowflake Documentation and cloud services (AWS, Azure, GCP) documentation on presigned URLs and stage interactions.

### NEW QUESTION: 300

Bulk data is loaded daily from an external storage location by using the following command:

```
COPY INTO source_table
FROM ' s3://source/data/files '
STORAGE_INTEGRATION = s3_source_int
FILE_FORMAT = (FORMAT_NAME = source_format)
PATTERN = ' sales.*[.]csv ' ;
```

How will the data be loaded?

- A. All files that match the specified format will be appended every time the command runs.
- B. Only new files will be loaded; previously loaded files will be skipped.
- C. All files matching s3://source/data/files/\*.csv will be loaded from an internal stage.
- D. The source table will be truncated, and all files from the location will be loaded.

**Answer:** [B \(LEAVE A REPLY\)](#)

Snowflake maintains file-loading metadata for each target table. When COPY INTO runs, Snowflake checks information such as the file name, size, checksum, and previous load status. Files that match the PATTERN parameter but have already been successfully loaded into the table are normally skipped, while newly discovered matching files are processed. This behavior helps prevent accidental duplicate loading when the same external location is scanned repeatedly. The command does not truncate the target table, and PATTERN limits the files considered without overriding load-history checks.

Previously loaded files can be loaded again by setting the FORCE copy option to TRUE, but doing so can create duplicate rows. Therefore, under the default behavior shown in the question, only new matching files are loaded.

### NEW QUESTION: 301

Which command will consume credits from a virtual warehouse?

- A. SELECT MAX(UUID) FROM SNOWFLAKE;
- B. SELECT COUNT(FLAKE\_ID) FROM SNOWFLAKE GROUP BY UUID;
- C. SHOW TABLES LIKE ' SNOWFL%' IN SCHEMA( ' PUBLIC ' );
- D. SELECT COUNT(\*) FROM SNOWFLAKE;

**Answer: A,B,D (LEAVE A REPLY)**

The correct answers are A, B, and D , assuming SNOWFLAKE is a table or queryable object.

Important exam-note correction:

This question is written as a single-answer question, but more than one listed command can consume virtual warehouse credits. Any SELECT query that scans or processes table data generally requires a running virtual warehouse and can consume warehouse credits.

Why A, B, and D are correct:

These are SELECT statements against a table-like object. Query execution that scans or aggregates data uses compute resources from a virtual warehouse.

Examples:

```
SELECT MAX(UUID) FROM SNOWFLAKE;  
SELECT COUNT(FLAKE_ID) FROM SNOWFLAKE GROUP BY UUID;  
SELECT COUNT(*) FROM SNOWFLAKE;
```

These involve SQL query processing and therefore require warehouse compute unless the result is fully served from cache.

Why C is incorrect:

SHOW commands are metadata commands handled by Snowflake's cloud services layer. They do not require an active virtual warehouse in the same way table-scanning SELECT queries do.

Official Snowflake documentation reference:

Snowflake documentation explains that virtual warehouses provide the compute resources required to execute SQL queries and DML operations. Metadata commands such as SHOW are handled by the cloud services layer.

Reference: Snowflake Documentation - Virtual warehouses; Snowflake Documentation - SHOW commands; SnowPro Core Study Guide - Snowflake Architecture and Resource Management.

==

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

### NEW QUESTION: 302

A company configures two resource monitors:

- \* An account-level resource monitor with a quota of 1,000 credits and the action NOTIFY & SUSPEND.

\* A warehouse-level resource monitor for the Marketing warehouse with a quota of 200 credits and the action NOTIFY & SUSPEND.

What happens when the Marketing warehouse reaches 200 credits while total account usage is only 700 credits?

- A. The Marketing warehouse is suspended immediately, and all running queries are canceled.
- B. A notification is sent, and the Marketing warehouse is suspended after currently running statements finish.
- C. A notification is sent, and the warehouse is suspended only when no ACCOUNTADMIN queries are running.
- D. The Marketing warehouse continues running until the account-level resource monitor reaches 1,000 credits.

**Answer:** (SHOW ANSWER)

A warehouse can be governed by its assigned warehouse-level resource monitor independently of an account-level resource monitor.

When the Marketing warehouse reaches its own 200-credit threshold, the NOTIFY & SUSPEND action takes effect even though the account-level monitor has not reached its 1,000-credit threshold.

The suspend action prevents new statements from starting and suspends the warehouse after currently running statements complete. It does not immediately cancel running queries.

Immediate cancellation corresponds to the NOTIFY & SUSPEND IMMEDIATE action. The user's administrative role does not exempt a running query from the monitor's configured threshold action.

References:

Snowflake Documentation - Resource monitors: <https://docs.snowflake.com/en/user-guide/resource-monitors> Snowflake SQL

Reference - CREATE RESOURCE MONITOR: <https://docs.snowflake.com/en/sql-reference/sql/create-resource-monitor>

### NEW QUESTION: 303

How does Snowflake support Multi-Factor Authentication (MFA) with a second factor device?

- A. Using time-based, one-time passcodes
- B. Through integration with OAuth
- C. Through Role-Based Access Control (RBAC)
- D. Using federated authentication and Single Sign-On (SSO)

**Answer:** A (LEAVE A REPLY)

### NEW QUESTION: 304

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

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

**Answer:** (SHOW ANSWER)

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: 305

Which automatic Snowpipe feature ensures that the same data file is not loaded multiple times into a table?

- A. Snowpipe deletes files from the stage after they are loaded, so they cannot be loaded twice.

- B.** Snowpipe references the metadata of previously loaded files and will not load new files with the same metadata.
- C.** Snowpipe renames each file after loading, ensuring that the same file cannot be loaded more than once.
- D.** Snowpipe moves each file to a temporary stage after loading, so it cannot be loaded more than once.

**Answer: (SHOW ANSWER)**

The correct answer is B. Snowpipe references the metadata of previously loaded files and will not load new files with the same metadata .

Snowflake keeps load metadata for files loaded into a table. Snowpipe uses this metadata to prevent duplicate loading of the same file into the same table.

Why B is correct:

Snowpipe tracks metadata for loaded files, including information that identifies files already processed. If the same file is encountered again, Snowpipe can use the load metadata to avoid loading it again.

Why the other options are incorrect:

- A). Snowpipe does not automatically delete files from the stage after loading.
- C). Snowpipe does not automatically rename staged files after loading.
- D). Snowpipe does not automatically move loaded files to a temporary stage.

Official Snowflake documentation reference:

Snowflake documentation explains that Snowpipe maintains load metadata and prevents the same files from being loaded multiple times into the same table by default.

Reference: Snowflake Documentation - Snowpipe; Snowflake Documentation - Load metadata; Snowflake Documentation - Data loading considerations; SnowPro Core Study Guide - Data Loading and Unloading.

## NEW QUESTION: 306

What Snowpark element represents a data set and provides ways to operate on the data?

- A.** List
- B.** Table
- C.** Array
- D.** DataFrame

**Answer: D (LEAVE A REPLY)**

The correct answer is D. DataFrame .

In Snowpark, a DataFrame represents a lazily evaluated relational data set. It provides methods to operate on data, such as selecting columns, filtering rows, joining data sets, grouping data, and applying transformations.

Why D is correct:

A Snowpark DataFrame is the primary abstraction for working with data in Snowpark. It represents data in Snowflake and lets developers build transformations using supported programming languages such as Python, Java, or Scala.

Example concept:

```
df = session.table( " CUSTOMERS " )
filtered_df = df.filter(df[ " COUNTRY " ] == " USA " )
```

The DataFrame represents the data set, and operations on it are translated into SQL that runs in Snowflake.

Why the other options are incorrect:

- A). A list is a general programming data structure, not the primary Snowpark data set abstraction.

B). A table stores data in Snowflake, but in Snowpark, the object used to represent and operate on a data set programmatically is a DataFrame.

C). An array is a data structure or semi-structured type, not the main Snowpark abstraction for data operations.

Official Snowflake documentation reference:

Snowflake documentation describes Snowpark DataFrames as objects that represent a lazily evaluated relational data set and provide methods for transforming and analyzing data.

Reference: Snowflake Documentation - Snowpark DataFrames; Snowflake Documentation - Snowpark API; SnowPro Core Study Guide - SQL and Snowflake Objects.

#### **NEW QUESTION: 307**

Which command will identify the names of files that are in a Snowflake stage?

**A.** SHOW FILE FORMAT

**B.** SELECT \$1 FROM < location >

**C.** LIST

**D.** DESCRIBE FILE FORMAT

**Answer:** ([SHOW ANSWER](#))

#### **NEW QUESTION: 308**

In the query profiler view for a query, which components represent areas that can be used to help optimize query performance? (Select TWO)

**A.** Bytes scanned

**B.** Bytes sent over the network

**C.** Number of partitions scanned

**D.** Percentage scanned from cache

**E.** External bytes scanned

**Answer:** **A,C** ([LEAVE A REPLY](#))

In the query profiler view, the components that represent areas that can be used to help optimize query performance include 'Bytes scanned' and 'Number of partitions scanned'. 'Bytes scanned' indicates the total amount of data the query had to read and is a direct indicator of the query's efficiency. Reducing the bytes scanned can lead to lower data transfer costs and faster query execution. 'Number of partitions scanned' reflects how well the data is clustered; fewer partitions scanned typically means better performance because the system can skip irrelevant data more effectively.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Query Profiling<sup>1</sup>

#### **NEW QUESTION: 309**

When an object is created in Snowflake, who owns the object?

**A.** The user 's default role

**B.** The owner of the parent schema

**C.** The public role

**D.** The current active primary role

**Answer: D ([LEAVE A REPLY](#))**

#### **NEW QUESTION: 310**

Which Snowflake edition offers the highest level of security for organizations that have the strictest requirements?

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

**Answer: D ([LEAVE A REPLY](#))**

The Virtual Private Snowflake (VPS) edition offers the highest level of security for organizations with the strictest security requirements. This edition provides a dedicated and isolated instance of Snowflake, including enhanced security features and compliance certifications to meet the needs of highly regulated industries or any organization requiring the utmost in data protection and privacy.

References:

Snowflake Documentation: Snowflake Editions

#### **NEW QUESTION: 311**

What does Snowflake recommend when planning virtual warehouse usage for a data load?

- A. Load the fewest possible number of large files.
- B. Use several single-duster warehouses.
- C. Dedicate a separate warehouse for loading data.
- D. Increase the size of the warehouse used.

**Answer: C ([LEAVE A REPLY](#))**

#### **NEW QUESTION: 312**

What information in a secure User-Defined Function, or UDF, is only available to authorized users? Select TWO.

- A. The UDF definition code
- B. Return type
- C. Parameter types
- D. List of data imports
- E. Handler language

**Answer: ([SHOW ANSWER](#))**

The correct answers are A. The UDF definition code and D. List of data imports .

A secure UDF is designed to protect sensitive implementation details. Secure UDFs help prevent unauthorized users from viewing the function body and other details that could expose business logic or sensitive references.

Why A is correct:

The definition code of a secure UDF is hidden from unauthorized users. This protects the logic used inside the function.

Why D is correct:

The list of imported files or data imports associated with the secure UDF is also protected from unauthorized users because it may reveal sensitive implementation details.

Why the other options are incorrect:

B). The return type is generally part of the function signature and must be visible so users can understand how to call and use the function.

C). Parameter types are also part of the function signature and are visible to users who need to call the UDF.

E). The handler language may be visible as part of function metadata and is not the best answer for protected secure UDF information.

Official Snowflake documentation reference:

Snowflake documentation explains that secure UDFs hide internal details, including the function definition and sensitive implementation information, from unauthorized users.

Reference: Snowflake Documentation - Secure UDFs; Snowflake Documentation - User-defined functions; SnowPro Core Study Guide - Security and Access Control.

### NEW QUESTION: 313

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

A. ACCOUNTADMIN

B. SECURITYADMIN

C. SYSADMIN

D. ORGADMIN

**Answer:** [\(SHOW ANSWER\)](#)

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: 314

What will be the result of this command?

```
SELECT value FROM LATERAL FLATTEN(input = > parse_json( ' [1, 2, 3] ' ));
```

A. It will return a single row with the array [1, 2, 3].

B. It will return a single row with three columns with values 1, 2, and 3.

C. It will return three rows with the values 1, 2, and 3.

D. It will return a single row with the values 1, 2, and 3 concatenated.

**Answer:** [B \(LEAVE A REPLY\)](#)

### NEW QUESTION: 315

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

A. True

B. False

**Answer:** [\(SHOW ANSWER\)](#)

The ACCOUNTADMIN role is the most powerful role in Snowflake and should be limited to a select number of users within an organization. It is responsible for account-level configurations and should not be used for day-to-day object creation or

management. Granting a custom role to ACCOUNTADMIN could inadvertently give broad access to users with this role, which is not a recommended security practice.

#### NEW QUESTION: 316

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

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

**Answer: B,E** ([LEAVE A REPLY](#))

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

References: Caching in the Snowflake Cloud Data Platform, Optimizing the warehouse cache

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 317

Which privilege must be granted by one role to another role, and cannot be revoked?

- A. MONITOR
- B. OPERATE
- C. OWNERSHIP
- D. ALL

**Answer: C** ([LEAVE A REPLY](#))

The OWNERSHIP privilege is unique in that it must be granted by one role to another and cannot be revoked.

This ensures that the transfer of ownership is deliberate and permanent, reflecting the importance of ownership in managing access and permissions.

#### NEW QUESTION: 318

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

- A. False
- B. True

**Answer: B** ([LEAVE A REPLY](#))

**NEW QUESTION: 319**

A Snowflake practitioner is tasked with designing a data lakehouse architecture that requires:

- \* Support for ACID transactions to ensure data consistency.
- \* The ability to evolve schemas over time without rewriting existing data.
- \* Efficient querying capabilities for large datasets.

Which table type meets these requirements?

- A.** Permanent
- B.** Iceberg
- C.** External
- D.** Transient

**Answer: B (LEAVE A REPLY)**

Apache Iceberg tables are designed for open data-lake and lakehouse architectures. The Iceberg table format supports ACID transactions, helping concurrent operations maintain atomicity, consistency, isolation, and durability. It also supports schema evolution, allowing table definitions to change without rewriting all existing data files. Iceberg metadata, snapshots, manifests, and hidden partitioning help compatible engines locate relevant data efficiently when querying large datasets. In Snowflake, Iceberg tables combine Snowflake query semantics and performance with data stored in externally managed cloud storage, usually in Parquet format. Permanent and transient tables are native Snowflake table types, while a basic external table primarily exposes staged files through metadata and is read-only.

Therefore, the Iceberg table type satisfies all three stated lakehouse requirements.

**NEW QUESTION: 320**

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 (LEAVE A REPLY)**

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: 321**

What happens when a Snowflake user changes the data retention period at the schema level?

- A.** All child objects will retain data for the new retention period.
- B.** All child objects that do not have an explicit retention period will automatically inherit the new retention period.
- C.** All child objects with an explicit retention period will be overridden with the new retention period.
- D.** All explicit child object retention periods will remain unchanged.

**Answer: B (LEAVE A REPLY)**

When the data retention period is changed at the schema level, all child objects that do not have an explicit retention period set will inherit the new retention period from the schema<sup>4</sup>.

**NEW QUESTION: 322**

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

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

**Answer:** ([SHOW ANSWER](#))

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

Snowflake Documentation: Access Control in Snowflake

**NEW QUESTION: 323**

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

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

**Answer:** A ([LEAVE A REPLY](#))

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

References:

Snowflake Documentation: Data Loading and Unloading

**NEW QUESTION: 324**

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

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

**Answer:** D ([LEAVE A REPLY](#))

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

References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 325**

Which roles are needed to access Document AI? Select TWO.

- A. PUBLIC
- B. SYSADMIN
- C. ACCOUNTADMIN
- D. OBJECT\_VIEWER
- E. DOCUMENT\_INTELLIGENCE\_CREATOR

**Answer: (SHOW ANSWER)**

The correct answers are C. ACCOUNTADMIN and E. DOCUMENT\_INTELLIGENCE\_CREATOR .

Document AI is part of Snowflake Cortex and uses specific privileges/roles for access and administration. The ACCOUNTADMIN role is commonly required for initial setup or granting required privileges, while DOCUMENT\_INTELLIGENCE\_CREATOR is used to create and work with Document AI objects.

Why C is correct:

ACCOUNTADMIN has the account-level authority needed to enable or manage required access for Document AI features.

Why E is correct:

DOCUMENT\_INTELLIGENCE\_CREATOR is the Snowflake role associated with creating Document AI- related objects and using Document AI capabilities.

Why the other options are incorrect:

A). PUBLIC is available to all users but does not provide Document AI creation privileges.

B). SYSADMIN manages many account objects, but it is not the specific Document AI role in the answer set.

D). OBJECT\_VIEWER is a Snowflake database role for object metadata visibility, not Document AI access.

Official Snowflake documentation reference:

Snowflake documentation for Document AI describes the required access model and includes the DOCUMENT\_INTELLIGENCE\_CREATOR role for creating Document AI objects, with account administrator involvement for setup.

Reference: Snowflake Documentation - Document AI; Snowflake Cortex; SnowPro Core Study Guide - Snowflake Architecture and Security.

## NEW QUESTION: 326

Which of the following are benefits of micro-partitioning? (Select TWO)

- A. Micro-partitions cannot overlap in their range of values
- B. Micro-partitions are immutable objects that support the use of Time Travel.
- C. Micro-partitions can reduce the amount of I/O from object storage to virtual warehouses
- D. Rows are automatically stored in sorted order within micro-partitions
- E. Micro-partitions can be defined on a schema-by-schema basis

**Answer: B,C (LEAVE A REPLY)**

Micro-partitions in Snowflake are immutable objects, which means once they are written, they cannot be modified. This immutability supports the use of Time Travel, allowing users to access historical data within a defined period. Additionally, micro-partitions can significantly reduce the amount of I/O from object storage to virtual warehouses. This is because Snowflake's query optimizer can skip over micro-partitions that do not contain relevant data for a query, thus reducing the amount of data that needs to be scanned and transferred.

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

<https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions.html>

### NEW QUESTION: 327

A query containing a WHERE clause is running longer than expected. The Query Profile shows that all micro-partitions are being scanned. How should this query be optimized?

- A. Create a view on the table.
- B. Add a clustering key to the table.
- C. Add a limit clause to the query.
- D. Add a Dynamic Data Masking policy to the table.

**Answer: B** ([LEAVE A REPLY](#))

When a query containing a WHERE clause is running longer than expected, and the Query Profile shows that all micro-partitions are being scanned, the query can be optimized by adding a clustering key to the table.

Understanding Micro-Partitioning in Snowflake:

Snowflake automatically partitions tables into micro-partitions for efficient storage and query performance.

Each micro-partition contains metadata about the range of values it holds, which helps in pruning irrelevant partitions during query execution.

Role of Clustering Keys:

A clustering key defines how data in a table is organized within micro-partitions.

By specifying a clustering key, you can control the physical layout of data, ensuring that related rows are stored together.

This organization improves query performance by reducing the number of micro-partitions that need to be scanned.

Optimizing Queries with Clustering Keys:

Adding a clustering key based on columns frequently used in WHERE clauses helps Snowflake quickly locate and scan relevant micro-partitions.

This minimizes the amount of data scanned and reduces query execution time.

Example:

```
ALTER TABLE my_table CLUSTER BY (column1, column2);
```

This command adds a clustering key to my\_table using column1 and column2.

Future queries that filter on these columns will benefit from improved performance.

Benefits:

Reduced query execution time: Fewer micro-partitions need to be scanned.

Improved resource utilization: More efficient data retrieval leads to lower compute costs.

Snowflake Documentation: Clustering Keys

Snowflake Documentation: Query Profile

### NEW QUESTION: 328

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

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

**Answer: (SHOW ANSWER)**

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

#### **NEW QUESTION: 329**

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

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

**Answer:** ([SHOW ANSWER](#))

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

References: Snowflake Documentation on Time Travel

#### **NEW QUESTION: 330**

Which Snowflake partner specializes in data catalog solutions?

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

**Answer:** ([SHOW ANSWER](#))

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: 331**

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:** C,D ([LEAVE A REPLY](#))

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

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 332

Which operator is used to evaluate query pruning efficiency by comparing partitions scanned to partitions total?

- A. Filter
- B. TableScan
- C. ExternalScan
- D. IndexScan

**Answer: B (LEAVE A REPLY)**

#### NEW QUESTION: 333

How can a Snowflake user validate data that is unloaded using the COPY INTO < location > command?

- A. Load the data into a CSV file.
- B. Load the data into a relational table.
- C. Use the VALIDATION\_MODE - SQL statement.
- D. Use the validation mode = return rows statement.

**Answer: C (LEAVE A REPLY)**

To validate data unloaded using the COPY INTO < location > command, a Snowflake user can use the VALIDATION\_MODE parameter within the SQL statement to test the files for errors without loading them<sup>3</sup>.

#### NEW QUESTION: 334

What will be the retention period for an object if the parameter MIN\_DATA\_RETENTION\_TIME\_IN\_DAYS is set at the account level and the DATA\_RETENTION\_TIME\_IN\_DAYS is set at the object level?

- A. The retention period will be equal to the lesser of the two values of parameters DATA\_RETENTION\_TIME\_IN\_DAYS set at the object level and MIN\_DATA\_RETENTION\_TIME\_IN\_DAYS set at the account level.
- B. The retention period will be equal to the value of the parameter DATA\_RETENTION\_TIME\_IN\_DAYS set at the object level.
- C. The retention period will be equal to the value of the parameter MIN\_DATA\_RETENTION\_TIME\_IN\_DAYS set at the account level.

**D.** The retention period will be equal to the greatest of the two values of parameters DATA\_RETENTION\_TIME\_IN\_DAYS set at the object level, and MIN\_DATA\_RETENTION\_TIME\_IN\_DAYS set at the account level.

**Answer: A** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 335**

When an ACCOUNTADMIN gives a user a custom role, what privilege or privileges is the user granted by default?

- A.** All privileges on objects allowed by the custom role
- B.** Access to all the objects owned by the USERADMIN role
- C.** Access to the PUBLIC role
- D.** All privileges that have been granted to the ACCOUNTADMIN

**Answer: C** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 336**

When should a user consider disabling auto-suspend for a virtual warehouse? (Select TWO).

- A.** When users will be using compute at different times throughout a 24/7 period
- B.** When managing a steady workload
- C.** When the compute must be available with no delay or lag time
- D.** When the user does not want to have to manually turn on the warehouse each time it is needed
- E.** When the warehouse is shared across different teams

**Answer: B,C** ([LEAVE A REPLY](#))

Disabling auto-suspend for a virtual warehouse is recommended when there is a steady workload, which ensures that compute resources are always available. Additionally, it is advisable to disable auto-suspend when immediate availability of compute resources is critical, eliminating any startup delay

#### **NEW QUESTION: 337**

Which roles are needed to access Document AI? Select TWO.

- A.** PUBLIC
- B.** SYSADMIN
- C.** ACCOUNTADMIN
- D.** OBJECT\_VIEWER
- E.** DOCUMENT\_INTELLIGENCE\_CREATOR

**Answer: C,E** ([LEAVE A REPLY](#))

The correct answers are C. ACCOUNTADMIN and E. DOCUMENT\_INTELLIGENCE\_CREATOR .

Document AI is part of Snowflake Cortex and uses specific privileges/roles for access and administration. The ACCOUNTADMIN role is commonly required for initial setup or granting required privileges, while DOCUMENT\_INTELLIGENCE\_CREATOR is used to create and work with Document AI objects.

Why C is correct:

ACCOUNTADMIN has the account-level authority needed to enable or manage required access for Document AI features.

Why E is correct:

DOCUMENT\_INTELLIGENCE\_CREATOR is the Snowflake role associated with creating Document AI- related objects and using Document AI capabilities.

Why the other options are incorrect:

- A). PUBLIC is available to all users but does not provide Document AI creation privileges.
- B). SYSADMIN manages many account objects, but it is not the specific Document AI role in the answer set.
- D). OBJECT\_VIEWER is a Snowflake database role for object metadata visibility, not Document AI access.

Official Snowflake documentation reference:

Snowflake documentation for Document AI describes the required access model and includes the DOCUMENT\_INTELLIGENCE\_CREATOR role for creating Document AI objects, with account administrator involvement for setup.

Reference: Snowflake Documentation - Document AI; Snowflake Cortex; SnowPro Core Study Guide - Snowflake Architecture and Security.

==

### NEW QUESTION: 338

A Snowflake user is querying JSON data and all values in the result set are wrapped in quotes. Which step will unwrap the data from all the columns in the query results?

- A. Use the TRIM function to remove the quotes.
- B. Use the INSTR function to remove the quotes.
- C. Use the SUBSTRING function to select the values that are inside the quotes.
- D. Use the CAST function to convert the columns to the correct data type.

**Answer: D** ([LEAVE A REPLY](#))

When querying VARIANT columns (such as JSON), the values are often returned as strings . To convert them to usable SQL types, users must apply the CAST() function (e.g., CAST(column:value AS NUMBER)).

Reference:

Snowflake Docs: Querying Semi-Structured Data

### NEW QUESTION: 339

A user creates a stage using the following command:

What will be the outcome?

- A. A stage with a directory table set to automatically refresh will be created.
- B. The command will fail to run because the name of the directory table is not specified.
- C. An error will be received stating that the storage location for the stage must be identified when creating a stage with a directory table.
- D. A stage with a directory table that has metadata that must be manually refreshed will be created.

**Answer: A** ([LEAVE A REPLY](#))

### NEW QUESTION: 340

Which Snowflake feature is used to track sensitive data to ensure compliance, identify data, and monitor resource usage?

- A. Data quality monitoring
- B. Data classification
- C. Multi-Factor Authentication (MFA)
- D. Object tagging

**Answer:** ([SHOW ANSWER](#))

**NEW QUESTION: 341**

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

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

**Answer: B** ([LEAVE A REPLY](#))

The `GET_PREIGNED_URL` function in Snowflake generates a pre-signed URL for a file in an internal stage.

This URL can be shared with external tools or users who do not have direct access to Snowflake, allowing them to download the file.

Generate Pre-Signed URL:

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

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

References:

Snowflake Documentation: `GET_PREIGNED_URL`

Snowflake Documentation: Working with Stages

**NEW QUESTION: 342**

What is a benefit of using system tags?

- A. To manage compute resources
- B. To support data governance
- C. To increase query execution speed
- D. To control data backup frequency

**Answer: B** ([LEAVE A REPLY](#))

System tags help enable data classification for governance, privacy, and compliance. For example, Snowflake uses system-defined tags like `SNOWFLAKE.CORE.DATA_CLASSIFICATION` to detect and label sensitive data (e.g., PII, financial info).

Reference:

Snowflake Docs: System Tags

**NEW QUESTION: 343**

Revoking which privilege will prevent one role from giving privileges to other roles?

- A. USAGE
- B. MODIFY
- C. REFERENCES
- D. MANAGE GRANTS

**Answer: D** ([LEAVE A REPLY](#))

The correct answer is D. MANAGE GRANTS .

The MANAGE GRANTS privilege allows a role to grant and revoke privileges on objects. If this privilege is revoked, the role loses the broad ability to manage privileges for other roles.

Why D is correct:

MANAGE GRANTS is a powerful global privilege. It allows a role to grant and revoke privileges, even on objects the role does not own, depending on Snowflake's access control rules.

Why the other options are incorrect:

A). USAGE allows access to an object or namespace, such as a database, schema, warehouse, or integration. It does not allow broad grant management.

B). MODIFY allows changes to certain object properties or data, depending on the object type, but it does not broadly allow granting privileges to roles.

C). REFERENCES is related to referencing objects, such as in constraints or policies. It does not control grant administration.

Official Snowflake documentation reference:

Snowflake documentation describes MANAGE GRANTS as the global privilege that enables granting and revoking privileges on objects.

Reference: Snowflake Documentation - Access control privileges; Snowflake Documentation - MANAGE GRANTS; SnowPro Core Study Guide - Security and Access Control.

### NEW QUESTION: 344

A Snowflake account administrator has set the resource monitors as shown in the diagram, with actions defined for each resource monitor as " Notify & Suspend Immediately " .

What is the MAXIMUM limit of credits that Warehouse 2 can consume?

A. 0

B. 1500

C. 3500

D. 5000

**Answer: B (LEAVE A REPLY)**

In the provided exhibit, Warehouse 2 is under Resource Monitor 2, which has a credit quota of 1000 credits.

However, it ' s important to note that warehouses can also draw from the account-level resource monitor if they consume their dedicated monitor ' s quota. In this exhibit, Resource Monitor 1 is set for the account with a credit quota of 500 credits.

The exhibit does not clearly show whether Warehouse 2 is also subject to the account-level Resource Monitor 1, but if we assume the typical case where all warehouses can draw from the account-level monitor, the maximum limit of credits Warehouse 2 can consume would be the sum of its dedicated resource monitor (Resource Monitor 2) and the account-level resource monitor (Resource Monitor 1).

Thus, Warehouse 2 can consume:

1000 credits from Resource Monitor 2

500 credits from Resource Monitor 1 (if applicable)

This gives a total of 1500 credits. Therefore, the maximum credit limit that Warehouse 2 can consume, taking into account the possibility of the account-level monitor ' s contribution, would be 1500 credits.

However, without explicit confirmation that Warehouse 2 can use credits from Resource Monitor 1, one would default to the specific resource monitor assigned to the warehouse, which is 1000 credits. Given the available options, the closest and most accurate answer is B. 1500, as it represents the sum of credits from both monitors that potentially apply to Warehouse 2.

References:

Snowflake Documentation on Resource Monitors: Managing Resource Monitors

### NEW QUESTION: 345

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:** [\(SHOW ANSWER\)](#)

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: 346

By default, which tasks can a user with the ORGADMIN role perform? Select TWO.

- A. Create a share.
- B. Enable replication.
- C. Create an account.
- D. Create a resource monitor.
- E. Create a reader account.

**Answer:** B,C [\(LEAVE A REPLY\)](#)

The correct answers are B. Enable replication and C. Create an account .

The ORGADMIN role is an organization-level role in Snowflake. It is used to manage organization-wide operations, including account management and enabling cross-account or cross-region capabilities such as replication.

Why B is correct:

ORGADMIN can enable replication for the organization. Replication is an organization-level capability because it can involve multiple accounts and regions.

Why C is correct:

ORGADMIN can create Snowflake accounts within the organization.

Why the other options are incorrect:

- A). Creating a share is an account-level data sharing task, not a default organization-level ORGADMIN task.
- D). Creating a resource monitor is an account-level resource management task, typically handled by account-level administrative roles such as ACCOUNTADMIN.
- E). Creating a reader account is an account-level secure data sharing task, not the primary default ORGADMIN responsibility in this answer set.

Official Snowflake documentation reference:

Snowflake documentation describes ORGADMIN as the organization administrator role used to manage organization-level operations, including creating accounts and enabling replication.

Reference: Snowflake Documentation - Organization administrator role; Snowflake Documentation - Creating accounts; Snowflake Documentation - Replication and failover; SnowPro Core Study Guide - Snowflake Account and Resource Management.

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 347

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

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

**Answer: A,C (LEAVE A REPLY)**

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

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

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

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

#### NEW QUESTION: 348

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

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

**Answer: D (LEAVE A REPLY)**

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: 349

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

A. BUILD\_SCOPED\_FILE\_URL

B. GET\_PREIGNED\_URL

C. BUILD\_STAGE\_FILE\_URL

D. GET\_STAGE\_LOCATION

Answer: ([SHOW ANSWER](#))

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

Generate Pre-Signed URL:

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

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

References:

Snowflake Documentation: GET\_PREIGNED\_URL

Snowflake Documentation: Working with Stages

#### NEW QUESTION: 350

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: ([SHOW ANSWER](#))

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: 351

When a replication schedule is set, how does Snowflake manage refreshes for failover groups?

A. Multiple refreshes will run concurrently.

B. Only one refresh will run at a time.

C. Refreshes will run regardless of role privileges.

D. If the current refresh was delayed, the next scheduled refresh will be canceled.

Answer: B ([LEAVE A REPLY](#))

The correct answer is B. Only one refresh will run at a time .

Failover groups can be refreshed on a schedule to replicate supported account objects and database objects to target accounts. Snowflake does not run overlapping refresh operations for the same failover group.

Why B is correct:

If a scheduled refresh is still running, Snowflake does not start another concurrent refresh for the same failover group. Refresh operations are managed so that only one refresh runs at a time.

Why the other options are incorrect:

A). Multiple refreshes for the same failover group do not run concurrently.

C). Refreshes still require the appropriate privileges and configured replication/failover group permissions.

D). A delayed or long-running refresh does not mean the next refresh is automatically canceled in the way stated. The key behavior is that refreshes do not overlap.

Official Snowflake documentation reference:

Snowflake documentation for replication and failover groups explains that scheduled refreshes are managed by Snowflake and that refresh operations do not run concurrently for the same group.

Reference: Snowflake Documentation - Replication and failover groups; Snowflake Documentation - Failover group refresh; SnowPro Core Study Guide - Snowflake Account and Resource Management.

### NEW QUESTION: 352

When is the VARIANT data type suitable when using semi-structured data? (Select TWO).

A. When the data set is huge and the order of the data elements in the data set does not change.

B. When the data contains nested data structures like JSON or XML.

C. When the data consists of images and videos.

D. When there is a need to add or modify elements within the data structure without interrupting the session.

E. When there is a need for a predefined schema which cannot be changed after it is created.

**Answer: B,D ([LEAVE A REPLY](#))**

### NEW QUESTION: 353

Taking which step will minimize the number of queries waiting to run on a virtual warehouse?

A. Increase the WAREHOUSE\_SIZE.

B. Increase the value of the AUTO\_SUSPEND parameter.

C. Set the WAIT\_FOR\_COMPLETION parameter to FALSE.

D. Increase the value of the MAX\_CLUSTER\_COUNT.

**Answer: D ([LEAVE A REPLY](#))**

The correct answer is D. Increase the value of the MAX\_CLUSTER\_COUNT .

Queries wait in a queue when a warehouse does not have enough available compute capacity to handle the current concurrency. For a multi-cluster warehouse, increasing MAX\_CLUSTER\_COUNT allows Snowflake to start more clusters when concurrency increases.

Why D is correct:

Increasing MAX\_CLUSTER\_COUNT allows a multi-cluster warehouse in Auto-scale mode to add more clusters, up to the configured maximum. This helps reduce query queuing for concurrent workloads.

Why the other options are incorrect:

- A). Increasing WAREHOUSE\_SIZE gives each cluster more compute power, which can help individual queries run faster, but it is not the best direct solution for reducing queued concurrent queries.
- B). Increasing AUTO\_SUSPEND keeps the warehouse running longer but does not increase concurrency capacity.
- C). WAIT\_FOR\_COMPLETION = FALSE is used with warehouse resize behavior and does not reduce query queuing.

Official Snowflake documentation reference:

Snowflake documentation explains that multi-cluster warehouses improve concurrency by starting additional clusters as query demand increases. The MAX\_CLUSTER\_COUNT property defines the maximum number of clusters that can run.

Reference: Snowflake Documentation - Multi-cluster warehouses; Snowflake Documentation - Warehouse properties; SnowPro Core Study Guide - Snowflake Account and Resource Management.

==

### NEW QUESTION: 354

How does Snowflake Copilot interact with customer data?

- A. Copilot analyzes customer data to optimize query performance.
- B. Copilot provides assistance when creating and interpreting queries.
- C. Copilot rewrites table schemas based on inferred patterns in customer datasets.
- D. Copilot executes stored procedures in response to natural language prompts.

**Answer: B (LEAVE A REPLY)**

The correct answer is B. Copilot provides assistance when creating and interpreting queries .

Snowflake Copilot is an AI-assisted feature that helps users work with Snowflake more efficiently. It can assist with writing SQL, understanding queries, and interpreting query-related information.

Why B is correct:

Copilot is designed to help users create, refine, and understand SQL queries using natural language assistance.

Why the other options are incorrect:

- A). Query optimization is primarily handled by Snowflake's optimizer and performance features, not Copilot.
- C). Copilot does not automatically rewrite table schemas based on inferred data patterns.
- D). Copilot is not primarily a feature for executing stored procedures from prompts.

Official Snowflake documentation reference:

Snowflake documentation describes Snowflake Copilot as an AI assistant for helping users generate, understand, and work with SQL and Snowflake data workflows.

Reference: Snowflake Documentation - Snowflake Copilot; Snowflake Cortex; SnowPro Core Study Guide - Snowflake Architecture.

==

### NEW QUESTION: 355

What does the NULL\_IF parameter do when loading CSV data into Snowflake?

- A. Converts binary data to text format.
- B. Converts specified strings to SQL NULL values.
- C. Removes NULL values from the dataset.
- D. Replaces every NULL value with an empty string.

**Answer: B (LEAVE A REPLY)**

NULL\_IF identifies one or more source-file strings that Snowflake should interpret as SQL NULL while loading data. For example, a file format can define NULL\_IF = ( ' NULL ' , ' NUL ' , ' \N ' , '' ).

When Snowflake encounters any specified value in the CSV input, it loads a SQL NULL into the corresponding target column. The parameter does not remove rows or values from the dataset, convert binary content, or automatically replace NULL values with empty strings. Snowflake applies the conversion to every occurrence of a configured value, regardless of whether that value visually resembles text or a number. When unloading data, NULL\_IF works in the opposite direction by converting SQL NULL values into the first specified output string, but this question specifically concerns loading CSV data.

### NEW QUESTION: 356

When regularly unloading similarly formatted semi-structured data, where should copy options be specified?

- A. In a table definition
- B. In a named file format
- C. In a named stage definition
- D. In the copy command

**Answer: B ([LEAVE A REPLY](#))**

The correct answer is B. In a named file format .

When data is regularly loaded or unloaded using the same format settings, Snowflake recommends defining those settings in a named file format. This avoids repeating the same file format options in every COPY command and reduces operational overhead.

Why B is correct:

A named file format stores reusable format options, such as file type and formatting behavior. For semi- structured data, named file formats can define options for formats such as JSON, Avro, ORC, Parquet, or XML.

Example:

```
CREATE FILE FORMAT my_json_format
TYPE = JSON
COMPRESSION = AUTO;
COPY INTO @my_stage/output/
FROM my_table
FILE_FORMAT = my_json_format;
```

Why the other options are incorrect:

- A). A table definition defines table columns and properties, not unload formatting options.
- C). A named stage can include a file format reference, but the reusable copy/file parsing settings themselves are defined in a named file format.
- D). Copy options can be specified directly in the COPY command, but that is less efficient for repeated, similarly formatted unload operations.

Official Snowflake documentation reference:

Snowflake documentation explains that named file formats are reusable database objects that define format options for loading and unloading data.

Reference: Snowflake Documentation - CREATE FILE FORMAT; Snowflake Documentation - Data unloading; Snowflake Documentation - COPY INTO < location > ; SnowPro Core Study Guide - Data Loading and Unloading.

**NEW QUESTION: 357**

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

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

**Answer: D,E** ([LEAVE A REPLY](#))

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

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Secure Data Sharing<sup>1</sup>

**NEW QUESTION: 358**

To use the overwrite option on insert, which privilege must be granted to the role?

- A. truncate
- B. DELETE
- C. UPDATE
- D. SELECT

**Answer: (**[SHOW ANSWER](#)**)**

To use the overwrite option on insert in Snowflake, the DELETE privilege must be granted to the role. This is because overwriting data during an insert operation implicitly involves deleting the existing data before inserting the new data.

Understanding the Overwrite Option: The overwrite option (INSERT OVERWRITE) allows you to replace existing data in a table with new data. This operation is particularly useful for batch-loading scenarios where the entire dataset needs to be refreshed.

Why DELETE Privilege is Required: Since the overwrite operation involves removing existing rows in the table, the executing role must have the DELETE privilege to carry out both the deletion of old data and the insertion of new data.

Granting DELETE Privilege:

To grant the DELETE privilege to a role, an account administrator can execute the following SQL command:

sqlCopy code

```
GRANT DELETE ON TABLE my_table TO ROLE my_role;
```

**NEW QUESTION: 359**

What objects can be cloned within Snowflake? (Select TWO).

- A. Schemas
- B. Users
- C. External tables
- D. Internal named stages
- E. External named stages

**Answer: (**[SHOW ANSWER](#)**)**

In Snowflake, cloning is available for certain types of objects, allowing quick duplication without copying data:

Schemas: These can be cloned, enabling users to replicate entire schema structures, including tables and views, for development or testing.

Internal named stages: These stages, used to store data files within Snowflake, can also be cloned, preserving configurations for data loading.

Users and external objects (like external stages or tables) cannot be cloned due to their dependency on external data and configurations outside Snowflake.

### NEW QUESTION: 360

What AUTHENTICATION\_METHODS value should be selected for the Single Sign-On, or SSO, login option to appear when creating a Snowflake authentication policy?

A. PASSWORD

B. KEYPAIR

C. SAML

D. OAUTH

**Answer: C (LEAVE A REPLY)**

The correct answer is C. SAML .

In Snowflake, Single Sign-On is commonly implemented using federated authentication through a SAML 2.0 identity provider. When creating an authentication policy, including SAML in the allowed AUTHENTICATION\_METHODS enables SSO-based login options.

Why C is correct:

SAML represents SAML-based federated authentication. This is the authentication method associated with SSO login in Snowflake.

Example concept:

```
CREATE AUTHENTICATION POLICY sso_policy  
AUTHENTICATION_METHODS = ( ' SAML ' );
```

Why the other options are incorrect:

A). PASSWORD allows username and password authentication, not SSO.

B). KEYPAIR allows key-pair authentication, commonly used for service accounts and applications.

D). OAUTH allows OAuth-based authentication, but the SSO login option in this context is associated with SAML.

Official Snowflake documentation reference:

Snowflake documentation describes authentication policies as controlling allowed authentication methods, including password, SAML, OAuth, and key-pair authentication. SAML is used for federated SSO authentication.

Reference: Snowflake Documentation - Authentication policies; Snowflake Documentation - Federated authentication and SSO; SnowPro Core Study Guide - Security and Access Control.

### NEW QUESTION: 361

Which pages are included in the Activity area of Snowsight? (Select TWO).

A. Contacts

B. Sharing settings

C. Copy History

D. Query History

E. Automatic Clustering History

**Answer: (SHOW ANSWER)**

The Activity area of Snowsight includes the Query History page, which allows users to monitor and view details about queries executed in their account, including performance data<sup>1</sup>. It also includes the Automatic Clustering History, which provides insights into the automatic clustering operations performed on tables<sup>2</sup>.

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

### NEW QUESTION: 362

In addition to the SELECT privilege, what other privilege must a role have to select from a table in database1.schema1?

A. The USAGE privilege on the table.

B. The OWNERSHIP privilege on the schema.

C. The REFERENCES privilege on the table and the schema.

D. The USAGE privilege on the database and the schema.

**Answer: D (LEAVE A REPLY)**

The correct answer is D. The USAGE privilege on the database and the schema .

To query a table in Snowflake, a role needs the SELECT privilege on the table. In addition, the role must have USAGE on the containing database and schema so it can resolve and access the object path.

Why D is correct:

For a table named:

database1.schema1.table1

A role generally needs:

USAGE ON DATABASE database1

USAGE ON SCHEMA database1.schema1

SELECT ON TABLE database1.schema1.table1

Without USAGE on the database and schema, the role cannot access the namespace that contains the table, even if SELECT exists on the table.

Why the other options are incorrect:

A). There is no USAGE privilege on a table for this purpose. SELECT is the table-level privilege needed to read data.

B). OWNERSHIP on the schema is not required to query a table.

C). REFERENCES is used for foreign key/reference-related behavior and is not required for simple SELECT queries.

Official Snowflake documentation reference:

Snowflake documentation explains that access to objects requires privileges on the object and appropriate USAGE privileges on parent containers, such as the database and schema.

Reference: Snowflake Documentation - Access control privileges; Snowflake Documentation - Overview of access control; SnowPro Core Study Guide - Security and Access Control.

**NEW QUESTION: 363**

What happens to deleted or modified records during the Time Travel retention period?

- A. They are permanently deleted.
- B. They are versioned for recovery.
- C. They are moved to a separate table.
- D. They are archived in a backup.

**Answer:** ([SHOW ANSWER](#))

**NEW QUESTION: 364**

What action should be taken if a Snowflake user wants to share a newly created object in a database with consumers?

- A. Use the automatic sharing feature for seamless access.
- B. Drop the object and then re-add it to the database to trigger sharing.
- C. Recreate the object with a different name in the database before sharing.
- D. Use the grant privilege ... TO share command to grant the necessary privileges.

**Answer:** D ([LEAVE A REPLY](#))

When a Snowflake user wants to share a newly created object in a database with consumers, the correct action to take is to use the `GRANT privilege ... TO SHARE` command to grant the necessary privileges for the object to be shared. This approach allows the object owner or a user with the appropriate privileges to share database objects such as tables, secure views, and streams with other Snowflake accounts by granting access to a named share.

The `GRANT` statement specifies which privileges are granted on the object to the share. The object remains in its original location; sharing does not duplicate or move the object. Instead, it allows the specified share to access the object according to the granted privileges.

For example, to share a table, the command would be:

```
GRANT SELECT ON TABLE new_table TO SHARE consumer_share;
```

This command grants the `SELECT` privilege on a table named `new_table` to a share named `consumer_share`, enabling the consumers of the share to query the table.

Automatic sharing, dropping and re-adding the object, or recreating the object with a different name are not required or recommended practices for sharing objects in Snowflake. The use of the `GRANT` statement to a share is the direct and intended method for this purpose.

Reference: Snowflake Documentation on Sharing Data (<https://docs.snowflake.com/en/user-guide/data-sharing-intro.html>)

**NEW QUESTION: 365**

Which column in the `ACCESS_HISTORY` view in Snowflake describes access history for a column and provides additional information on how columns have been altered over time?

- A. `object_modified_by_dd1`
- B. `objects_modified`
- C. `base_objects_accessed`
- D. `direct_objects_accessed`

**Answer:** ([SHOW ANSWER](#))

**NEW QUESTION: 366**

What compute resource is used when loading data using Snowpipe?

- A. Snowpipe uses virtual warehouses provided by the user.
- B. Snowpipe uses an Apache Kafka server for its compute resources.
- C. Snowpipe uses compute resources provided by Snowflake.
- D. Snowpipe uses cloud platform compute resources provided by the user.

**Answer:** C ([LEAVE A REPLY](#))

Snowpipe is Snowflake's continuous data ingestion service that allows for loading data as soon as it's available in a cloud storage stage. Snowpipe uses compute resources managed by Snowflake, separate from the virtual warehouses that users create for querying data. This means that Snowpipe operations do not consume the computational credits of user-created virtual warehouses, offering an efficient and cost-effective way to continuously load data into Snowflake.

References:

Snowflake Documentation: Understanding Snowpipe

**NEW QUESTION: 367**

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 ([LEAVE A REPLY](#))

When using the ALLOW\_CLIENT\_MFA\_CACHING parameter, a cached Multi-Factor Authentication (MFA) token is valid for up to 4 hours. This allows for continuous, secure connectivity without users needing to respond to an MFA prompt at the start of each connection attempt to Snowflake within this timeframe.

**NEW QUESTION: 368**

In which hierarchy is tag inheritance possible?

- A. Organization - Account - Role
- B. Account - User - Schema
- C. Database - View - Column
- D. Schema - Table - Column

**Answer:** ([SHOW ANSWER](#))

In Snowflake, tag inheritance is a feature that allows tags, which are key-value pairs assigned to objects for the purpose of data governance and metadata management, to be inherited within a hierarchy. The hierarchy in which tag inheritance is possible is from Schema to Table to Column. This means that a tag applied to a schema can be inherited by the tables within that schema, and a tag applied to a table can be inherited by the columns within that table. References: Snowflake Documentation on Tagging and Object Hierarchy

### NEW QUESTION: 369

A query containing a WHERE clause is running longer than expected. The Query Profile shows that all micro-partitions are being scanned. How should this query be optimized?

- A. Create a view on the table.
- B. Add a clustering key to the table.
- C. Add a limit clause to the query.
- D. Add a Dynamic Data Masking policy to the table.

**Answer: B** ([LEAVE A REPLY](#))

When a query containing a WHERE clause is running longer than expected, and the Query Profile shows that all micro-partitions are being scanned, the query can be optimized by adding a clustering key to the table.

Understanding Micro-Partitioning in Snowflake:

Snowflake automatically partitions tables into micro-partitions for efficient storage and query performance.

Each micro-partition contains metadata about the range of values it holds, which helps in pruning irrelevant partitions during query execution.

Role of Clustering Keys:

A clustering key defines how data in a table is organized within micro-partitions.

By specifying a clustering key, you can control the physical layout of data, ensuring that related rows are stored together.

This organization improves query performance by reducing the number of micro-partitions that need to be scanned.

Optimizing Queries with Clustering Keys:

Adding a clustering key based on columns frequently used in WHERE clauses helps Snowflake quickly locate and scan relevant micro-partitions.

This minimizes the amount of data scanned and reduces query execution time.

Example:

```
ALTER TABLE my_table CLUSTER BY (column1, column2);
```

This command adds a clustering key to my\_table using column1 and column2.

Future queries that filter on these columns will benefit from improved performance.

Benefits:

Reduced query execution time: Fewer micro-partitions need to be scanned.

Improved resource utilization: More efficient data retrieval leads to lower compute costs.

References:

Snowflake Documentation: Clustering Keys

Snowflake Documentation: Query Profile

### NEW QUESTION: 370

Which categories are included in the execution time summary in a Query Profile? (Select TWO).

- A. Pruning
- B. Spilling
- C. Initialization
- D. Local Disk I/O
- E. Percentage of data read from cache

**Answer: A,C** ([LEAVE A REPLY](#))

In the execution time summary of a Query Profile in Snowflake, the categories included provide insights into various aspects of query execution. " Pruning " refers to the process by which Snowflake reduces the amount of data scanned by eliminating partitions of data that are not relevant to the query, thus improving performance. " Initialization " represents the time taken for query planning and setup before actual execution begins. These metrics are crucial for understanding and optimizing query performance.

Reference: Snowflake Documentation on the Query Profile, which outlines the different metrics and categories included in the execution summary: <https://docs.snowflake.com/en/user-guide/ui-query-profile.html#execution-summary>

### NEW QUESTION: 371

What costs are associated with sharing Snowflake data with a reader account ?

- A. The data provider incurs costs for the data storage costs, and the data consumer incurs compute costs.
- B. The data provider incurs costs for compute, and the data consumer incurs costs for data storage.
- C. The data provider does not incur costs when sharing data; the data consumer incurs costs for data storage and compute.
- D. The data consumer does not incur costs when accessing shared data; the data provider incurs costs for data storage and compute.

**Answer:** (SHOW ANSWER)

With reader accounts , the provider pays for:

- \* Storage of shared data
- \* Compute used by the reader account

The consumer does not pay , as the infrastructure is provisioned and managed by the provider.

Reference:

Snowflake Docs: Reader Accounts and Costs

### NEW QUESTION: 372

What is the SNOWFLAKE.ACCOUNT\_USAGE view that contains information about which objects were read by queries within the last 365 days (1 year)?

- A. VIEWS\_HISTORY
- B. OBJECT\_HISTORY
- C. ACCESS\_HISTORY
- D. LOGIN\_HISTORY

**Answer:** C (LEAVE A REPLY)

The ACCESS\_HISTORY view in the SNOWFLAKE.ACCOUNT\_USAGE schema contains information about the access history of Snowflake objects, such as tables and views, within the last 365 days<sup>1</sup>.

### NEW QUESTION: 373

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

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

**D. Tagging**

**Answer: C (LEAVE A REPLY)**

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

References:

Snowflake Documentation: Search Optimization Service

#### **NEW QUESTION: 374**

How do Snowflake data providers share data that resides in different databases?

**A. External tables**

**B. Secure views**

**C. Materialized views**

**D. User-Defined Functions (UDFs)**

**Answer: B (LEAVE A REPLY)**

Snowflake data providers can share data residing in different databases through secure views. Secure views allow for the referencing of objects such as schemas, tables, and other views contained in one or more databases, as long as those databases belong to the same account. This enables providers to share data securely and efficiently with consumers.

References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

#### **NEW QUESTION: 375**

A large enterprise requires a disaster recovery solution that ensures all critical databases and account-level objects can be quickly redirected to a secondary region hosted on a different cloud provider. Which solution will enable synchronized cross-cloud and cross-region replication to support this level of redundancy?

**A. Snowgrid**

**B. Data Clean Rooms**

**C. Cloning to an external stage**

**D. Replication and failover groups**

**Answer: D (LEAVE A REPLY)**

Replication and failover groups are used to replicate databases and supported account-level objects across accounts, regions, and cloud platforms. Failover groups support disaster recovery by allowing clients and workloads to be redirected to a secondary account.

Reference: Snowflake Documentation - Replication and failover groups; Business continuity and disaster recovery.

#### **NEW QUESTION: 376**

What is the least-privileged database role needed to view the definition of a secure materialized view?

**A. USAGE\_VIEWER**

**B. OBJECT\_VIEWER**

**C. SECURITY\_VIEWER**

**D. GOVERNANCE\_VIEWER**

**Answer: A (LEAVE A REPLY)**

The USAGE\_VIEWER role is the minimum required to view metadata and definitions of governance- enabled or secure objects like secure materialized views .

Reference:

Snowflake Docs: Governance Roles

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### **NEW QUESTION: 377**

Which Snowflake layer is associated with virtual warehouses?

- A. Cloud services
- B. Query processing
- C. Elastic memory
- D. Database storage

**Answer: B (LEAVE A REPLY)**

The layer of Snowflake ' s architecture associated with virtual warehouses is the Query Processing layer.

Virtual warehouses in Snowflake are dedicated compute clusters that execute SQL queries against the stored data. This layer is responsible for the entire query execution process, including parsing, optimization, and the actual computation. It operates independently of the storage layer, enabling Snowflake to scale compute and storage resources separately for efficiency and cost-effectiveness.

References:

Snowflake Documentation: Snowflake Architecture

#### **NEW QUESTION: 378**

What is required to set a stored procedure to EXECUTE AS? (Select TWO).

- A. Owner ' s rights
- B. ACCOUNTADMIN role
- C. SCHEMA\_USER
- D. Caller ' s rights
- E. USER

**Answer: (SHOW ANSWER)**

#### **NEW QUESTION: 379**

Which privilege must be granted by one role to another role, and cannot be revoked?

- A. MONITOR
- B. OPERATE

**C. OWNERSHIP**

**D. ALL**

**Answer: C ([LEAVE A REPLY](#))**

The OWNERSHIP privilege is unique in that it must be granted by one role to another and cannot be revoked. This ensures that the transfer of ownership is deliberate and permanent, reflecting the importance of ownership in managing access and permissions.

#### **NEW QUESTION: 380**

Which commands should be used to grant the privilege allowing a role to select data from all current tables and any tables that will be created later in a schema? (Choose two.)

**A.** grant USAGE on all tables in schema DB1.SCHEMA to role MYROLE;

**B.** grant USAGE on future tables in schema DB1.SCHEMA to role MYROLE;

**C.** grant SELECT on all tables in schema DB1.SCHEMA to role MYROLE;

**D.** grant SELECT on future tables in schema DB1.SCHEMA to role MYROLE;

**E.** grant SELECT on all tables in database DB1 to role MYROLE;

**F.** grant SELECT on future tables in database DB1 to role MYROLE;

**Answer: C,D ([LEAVE A REPLY](#))**

To grant a role the privilege to select data from all current and future tables in a schema, two separate commands are needed. The first command grants the SELECT privilege on all existing tables within the schema, and the second command grants the SELECT privilege on all tables that will be created in the future within the same schema.

#### **NEW QUESTION: 381**

Which columns are part of the result set of the Snowflake LATERAL FLATTEN command? (Choose two.)

**A.** CONTENT

**B.** PATH

**C.** BYTE\_SIZE

**D.** INDEX

**E.** DATATYPE

**Answer: ([SHOW ANSWER](#))**

The LATERAL FLATTEN command in Snowflake produces a result set that includes several columns, among which PATH and INDEX are included. PATH indicates the path to the element within a data structure that needs to be flattened, and INDEX represents the index of the element if it is an array<sup>2</sup>.

#### **NEW QUESTION: 382**

A column named " Data " contains VARIANT data and stores values as follows:

How will Snowflake extract the employee ' s name from the column data?

**A.** Data:employee.name

**B.** DATA:employee.name

**C.** data:Employee.name

**D.** data:employee.name

**Answer: ([SHOW ANSWER](#))**

In Snowflake, to extract a specific value from a VARIANT column, you use the column name followed by a colon and then the key. The keys are case-sensitive. Therefore, to extract the employee's name from the "Data" column, the correct syntax is data:employee.name.

#### **NEW QUESTION: 383**

What parameter setting will maximize parallel processing when unloading data from Snowflake?

- A. SINGLE = TRUE
- B. SINGLE = FALSE
- C. DETAILED\_OUTPUT
- D. DETAILED OUTPUT = TRUE/FALSE

**Answer: B** ([LEAVE A REPLY](#))

When SINGLE = FALSE is set during COPY INTO, Snowflake writes multiple output files , enabling parallelism . Setting SINGLE = TRUE forces one output file, reducing parallel processing.

Reference:

Snowflake Docs: COPY INTO Parameters

#### **NEW QUESTION: 384**

Which Snowflake edition enables data sharing only through Snowflake Support?

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

**Answer: A** ([LEAVE A REPLY](#))

The Snowflake edition that enables data sharing only through Snowflake Support is the Virtual Private Snowflake (VPS).By default, VPS does not permit data sharing outside of the VPS environment, but it can be enabled through Snowflake Support4.

#### **NEW QUESTION: 385**

A user needs to create a materialized view in the schema MYDB.MYSCHEMA.

Which statements will provide this access?

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

**Answer: D** ([LEAVE A REPLY](#))

In Snowflake, to create a materialized view, the user must have the necessary privileges on the schema where the view will be created. These privileges are granted through roles, not directly to individual users.

Therefore, the correct process is to grant the role to the user and then grant the privilege to create the materialized view to the role itself.

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

References:

Snowflake Documentation on Roles

Snowflake Documentation on Materialized Views

#### **NEW QUESTION: 386**

What will happen to partitions if a table is cloned and the original table type is then changed from permanent to transient?

- A.** Both old and new partitions will be changed to transient, and the original table type will change to transient.
- B.** Old partitions will be changed to transient, and any new partitions added to the clone will follow the transient lifecycle.
- C.** Old partitions will not be changed to transient, but any new partitions added to the clone will follow the transient lifecycle.
- D.** Both old and new partitions will be changed to transient, but the original table will keep the original table type.

**Answer: C** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 387**

What is the relationship between a Query Profile and a virtual warehouse?

- A.** A Query Profile can help users right-size virtual warehouses.
- B.** A Query Profile defines the hardware specifications of the virtual warehouse.
- C.** A Query Profile can help determine the number of virtual warehouses available.
- D.** A Query Profile automatically scales the virtual warehouse based on the query complexity.

**Answer: A** ([LEAVE A REPLY](#))

A Query Profile provides detailed execution information for a query, which can be used to analyze the performance and behavior of queries. This information can help users optimize and right-size their virtual warehouses for better efficiency.

References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

#### **NEW QUESTION: 388**

Which SQL statement will require a virtual warehouse to run?

- A.** `SELECT COUNT(*) FROM TBL_EMPLOYEE;`
- B.** `ALTER TABLE TBL_EMPLOYEE ADD COLUMN EMP_REGION VARCHAR(20);`
- C.** `INSERT INTO TBL_EMPLOYEE(EMP_ID, EMP_NAME, EMP_SALARY, DEPT) VALUES(1,'Adam',20000,'Finance');`
- D.** `CREATE OR REPLACE TABLE TBL_EMPLOYEE (`

**Answer: (SHOW ANSWER)**

A virtual warehouse in Snowflake is used to perform data processing tasks that require computational resources, such as queries that modify data or perform significant computation. Of the options provided:

C . `INSERT INTO TBL_EMPLOYEE(EMP_ID, EMP_NAME, EMP_SALARY, DEPT) VALUES(1,'Adam',20000,'Finance');` This SQL statement performs a data modification operation (DML) by inserting a new record into the TBL\_EMPLOYEE table, which requires computational resources provided by a virtual warehouse to execute.

References:

Snowflake Documentation: Understanding Virtual Warehouses

### NEW QUESTION: 389

While running a query on a virtual warehouse in auto-scale mode, additional clusters are stated immediately if which setting is configured?

- A. MAX\_CLUSTER\_COUNT is increased and new\_max\_clusters is greater than running\_clusters
- B. MIN\_CLUSTER\_COUNT is decreased and new\_min\_clusters is less than running\_clusters
- C. MAX\_CLUSTER\_COUNT is decreased and new\_max\_clusters is less than running\_clusters
- D. MIN\_CLUSTER\_COUNT is increased and new\_min\_clusters is greater than running\_clusters

**Answer: D** ([LEAVE A REPLY](#))

### NEW QUESTION: 390

A user frequently runs a reporting query that returns the total record count and the minimum and maximum timestamp from a 100 TB fact table.

The table is loaded once per week and remains static between loads. The reporting query runs hundreds of times per day and causes significant warehouse-credit consumption because it repeatedly scans the fact table.

Which object will improve performance, minimize repeated query-compute costs, and provide the fastest response time?

- A. Standard view
- B. Dynamic table
- C. Temporary table
- D. Materialized view

**Answer: D** ([LEAVE A REPLY](#))

A materialized view stores precomputed query results. Snowflake maintains the materialized view when its underlying base-table data changes. Because the source table changes only once per week, maintenance occurs much less frequently than the reporting query is executed.

Repeated queries can read the much smaller precomputed result instead of scanning the 100 TB base table hundreds of times per day. Snowflake's optimizer can also automatically rewrite an eligible query to use a materialized view even when the query references the base table.

A standard view stores only its SQL definition and normally executes the underlying query whenever it is referenced. A temporary table would require manual creation, refresh, and lifecycle management. A dynamic table could materialize transformed results, but it requires refresh scheduling and refresh compute. For this repeated, stable aggregation pattern, the intended optimization object is a materialized view.

Materialized-view maintenance itself can consume credits, so it is most beneficial when query savings outweigh maintenance costs. The table's infrequent load schedule makes that condition more likely.

References:

Snowflake Documentation - Working with materialized views: <https://docs.snowflake.com/en/user-guide/views-materialized>

Snowflake Documentation - Deciding when to create a materialized view: <https://docs.snowflake.com/en/user-guide/views-materialized#deciding-when-to-create-a-materialized-view>

==

#### NEW QUESTION: 391

What is included in the calculation of the execution time reported in a Query Profile? (Select TWO).

- A. Spilling
- B. Local disk IO
- C. Caching
- D. Processing
- E. Pruning

Answer: A,D ([LEAVE A REPLY](#))

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 392

How can storage costs be MINIMIZED while maintaining the history and integrity of a large, permanent table that is updated frequently?

- A. Use zero-copy clones of the table.
- B. Convert the table type to transient.
- C. Enable automatic clustering
- D. Decrease the Time Travel retention period

Answer: D ([LEAVE A REPLY](#))

#### NEW QUESTION: 393

How does the search optimization service help Snowflake users improve query performance?

- A. It scans the micro-partitions based on the joins used in the queries and scans only join columns.
- B. It maintains a persistent data structure that keeps track of the values of the table's columns in each of its micro-partitions.
- C. It scans the local disk cache to avoid scans on the tables used in the Query.
- D. It keeps track of running queries and their results and saves those extra scans on the table.

Answer: B ([LEAVE A REPLY](#))

The search optimization service in Snowflake enhances query performance by maintaining a persistent data structure. This structure indexes the values of table columns across micro-partitions, allowing Snowflake to quickly identify which micro-partitions contain relevant data for a query. By efficiently narrowing down the search space, this service reduces the amount of data scanned during query execution, leading to faster response times and more efficient use of resources. References: Snowflake Documentation on Search Optimization Service

#### NEW QUESTION: 394

Which REST API can be used with unstructured data?

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

**Answer: C** ([LEAVE A REPLY](#))

The REST API used with unstructured data in Snowflake is GET /api/files/, which retrieves (downloads) a data file from an internal or external stage<sup>4</sup>.

#### NEW QUESTION: 395

Which command is used to unload files from an internal or external stage to a local file system?

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

**Answer: B** ([LEAVE A REPLY](#))

The command used to unload files from an internal or external stage to a local file system in Snowflake is the GET command. This command allows users to download data files that have been staged, making them available on the local file system for further use<sup>23</sup>.

#### NEW QUESTION: 396

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** ([LEAVE A REPLY](#))

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: 397

True or False: Loading data into Snowflake requires that source data files be no larger than 16MB.

- A. True
- B. False

**Answer: (SHOW ANSWER)**

Snowflake does not require source data files to be no larger than 16MB. In fact, Snowflake recommends that for optimal load performance, data files should be roughly 100-250 MB in size when compressed. However, it is not recommended to load very large files (e.g., 100 GB or larger) due to potential delays and wasted credits if errors occur. Smaller files should be aggregated to minimize processing overhead, and larger files should be split to distribute the load among compute resources in an active warehouse.

References: Preparing your data files | Snowflake Documentation

**NEW QUESTION: 398**

A user needs to connect a private GitHub repository to Snowflake for stored-procedure version control. The user's organization requires secure authentication.

Which authentication method should be configured for access to the private GitHub repository?

- A. Username and account password
- B. Personal access token
- C. OAuth token
- D. Multi-factor authentication

**Answer: (SHOW ANSWER)**

For a private GitHub repository accessed over HTTPS, Snowflake can authenticate using a GitHub personal access token. The token is stored securely in a Snowflake secret rather than written directly into the Git repository object.

The secret is referenced by the API integration or Git repository configuration. For GitHub, a personal access token is generally supplied in the password field of a password-type Snowflake secret.

A normal GitHub account password is not the recommended authentication mechanism for Git operations.

Interactive MFA cannot be performed by Snowflake when it synchronizes a repository. An arbitrary OAuth token is not the standard answer for the documented private-GitHub setup represented by this question.

References:

Snowflake Documentation - Setting up Snowflake to use Git: <https://docs.snowflake.com/en/developer-guide/git/git-setting-up>

Snowflake Documentation - Using a Git repository in Snowflake: <https://docs.snowflake.com/en/developer-guide/git/git-overview>

==

**NEW QUESTION: 399**

What unit is used to calculate the monthly cost for storing data in Snowflake?

- A. Gigabyte (GB)
- B. Megabyte (MB)
- C. Terabyte (TB)
- D. Petabyte (PB)

**Answer: (SHOW ANSWER)**

Snowflake bills data storage monthly in units of compressed terabytes (TB) . Usage below 1 TB is prorated.

Reference:

Snowflake Docs: Storage Costs

**NEW QUESTION: 400**

How does the authorization associated with a pre-signed URL work for an unstructured file?

- A. The role specified in the GET REST API call must have sufficient privileges on the stage to access the referenced file using the URL.
- B. Only the users who have roles with sufficient privileges on the URL can access the referenced file.
- C. Only the user who generates the URL can use the URL to access the referenced file.

D. Anyone who has the URL can access the referenced file for the life of the token.

**Answer: A** ([LEAVE A REPLY](#))

#### NEW QUESTION: 401

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

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

**Answer: (SHOW ANSWER)**

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

#### NEW QUESTION: 402

Which Query Profile metrics can be used to determine if a query is pruning efficiently?

- A. Number of rows inserted and Number of rows updated
- B. Partitions scanned and Partitions total
- C. Bytes read from result and Bytes written to result
- D. Bytes spilled to local storage and Bytes spilled to remote storage

**Answer: B** ([LEAVE A REPLY](#))

#### NEW QUESTION: 403

Which table function is used to perform additional processing on the results of a previously-run query?

- A. QUERY\_HISTORY
- B. RESULT\_SCAN
- C. DESCRIBE\_RESULTS
- D. QUERY HISTORY BY SESSION

**Answer: B** ([LEAVE A REPLY](#))

The `RESULT_SCAN` table function is used in Snowflake to perform additional processing on the results of a previously-run query. It allows users to reference the result set of a previous query by its query ID, enabling further analysis or transformations without re-executing the original query.

Snowflake Documentation: `RESULT_SCAN`

#### NEW QUESTION: 404

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 Aggregate Operator node is present
- D. The query is spilling to remote storage

**Answer: (SHOW ANSWER)**

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-C03 ] SnowPro Core Certification Exam Study Guide

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

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

### NEW QUESTION: 405

What is the purpose of key-pair authentication?

- A. To encrypt data that can be accessed with one of two keys
- B. To allow multiple users to securely share data
- C. To securely authenticate a user or an application
- D. To enforce Multi-Factor Authentication, or MFA

**Answer: C** ([LEAVE A REPLY](#))

The correct answer is C. To securely authenticate a user or an application .

Key-pair authentication uses a public/private key pair to authenticate to Snowflake. It is commonly used for service accounts, automation, applications, and client tools that need secure authentication without relying on passwords.

Why C is correct:

In key-pair authentication, the private key is used by the client, and the public key is associated with the Snowflake user. Snowflake verifies the authentication request using the public key.

Why the other options are incorrect:

- A). Key-pair authentication is about authentication, not encrypting Snowflake data for access by one of two keys.
- B). Secure data sharing is a separate Snowflake feature and is not the purpose of key-pair authentication.
- D). MFA is a separate authentication control. Key-pair authentication does not itself enforce MFA.

Official Snowflake documentation reference:

Snowflake documentation describes key-pair authentication as an authentication method using a public /private key pair for users and applications.

Reference: Snowflake Documentation - Key-pair authentication; Snowflake Documentation - User authentication; SnowPro Core Study Guide - Security and Access Control.

### NEW QUESTION: 406

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

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

**Answer: (**[SHOW ANSWER](#)**)**

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

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 407

Where can file formats be specified when unloading data from Snowflake? Select TWO.

- A. In a SELECT statement
- B. In a CREATE STAGE statement
- C. In a CREATE DATABASE statement
- D. In a CREATE INTEGRATION statement
- E. In a COPY INTO < location > statement

**Answer: B,E (LEAVE A REPLY)**

The correct answers are B. In a CREATE STAGE statement and E. In a COPY INTO < location > statement .

When unloading data from Snowflake, file format options can be specified directly in the COPY INTO < location > command or associated with a named stage. Defining the file format in a stage is useful when repeated loads or unloads use the same format settings.

Why B is correct:

A stage can include a file format definition or reference. When data is unloaded to that stage, the file format associated with the stage can be used.

Example:

```
CREATE STAGE my_stage
FILE_FORMAT = (TYPE = CSV FIELD_DELIMITER = ' , ' );
COPY INTO @my_stage
FROM my_table;
```

Why E is correct:

The COPY INTO < location > command can specify file format options directly.

Example:

```
COPY INTO @my_stage/output/
FROM my_table
FILE_FORMAT = (TYPE = CSV COMPRESSION = GZIP);
```

Why the other options are incorrect:

A). A SELECT statement defines the query result to unload, but file format options are not specified as part of a normal SELECT.

C). A database definition does not define unload file format options.

D). An integration defines access to external services or storage, not the file format used for unloading.

Official Snowflake documentation reference:

Snowflake documentation explains that file format options can be specified in the COPY INTO < location > command or as part of a named stage definition.

Reference: Snowflake Documentation - COPY INTO < location > ; Snowflake Documentation - CREATE STAGE; Snowflake Documentation - File formats; SnowPro Core Study Guide - Data Loading and Unloading.

==

### NEW QUESTION: 408

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: (SHOW ANSWER)**

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: 409

What will happen if a Snowflake user increases the size of a suspended virtual warehouse?

A. The provisioning of new compute resources for the warehouse will begin immediately.

B. The warehouse will remain suspended but new resources will be added to the query acceleration service.

C. The provisioning of additional compute resources will be in effect when the warehouse is next resumed.

D. The warehouse will resume immediately and start to share the compute load with other running virtual warehouses.

**Answer: C (LEAVE A REPLY)**

When a Snowflake user increases the size of a suspended virtual warehouse, the changes to compute resources are queued but do not take immediate effect. The provisioning of additional compute resources occurs only when the warehouse is resumed. This ensures that resources are allocated efficiently, aligning with Snowflake's commitment to cost-effective and on-demand scalability.

References:

**NEW QUESTION: 410**

In an auto-scaling multi-cluster virtual warehouse with the setting SCALING\_POLICY = ECONOMY enabled, when is another cluster started?

- A. When the system has enough load for 2 minutes
- B. When the system has enough load for 6 minutes
- C. When the system has enough load for 8 minutes
- D. When the system has enough load for 10 minutes

**Answer: A** ([LEAVE A REPLY](#))

In an auto-scaling multi-cluster virtual warehouse with the SCALING\_POLICY set to ECONOMY, another cluster is started when the system has enough load for 2 minutes (A). This policy is designed to optimize the balance between performance and cost, starting additional clusters only when the sustained load justifies it<sup>2</sup>.

**NEW QUESTION: 411**

Which types of charts does Snowsight support? (Select TWO).

- A. Area charts
- B. Bar charts
- C. Column charts
- D. Radar charts
- E. Scorecards

**Answer: A,B** ([LEAVE A REPLY](#))

Snowsight, Snowflake's user interface for executing and analyzing queries, supports various types of visualizations to help users understand their data better. Among the supported types, area charts and bar charts are two common options. Area charts are useful for representing quantities through the use of filled areas on the graph, often useful for showing volume changes over time. Bar charts, on the other hand, are versatile for comparing different groups or categories of data. Both chart types are integral to data analysis, enabling users to visualize trends, patterns, and differences in their data effectively.

References: Snowflake Documentation on Snowsight Visualizations

**NEW QUESTION: 412**

In Snowpipe, what is the default behavior of the ON\_ERROR copy option?

- A. CONTINUE
- B. SKIP\_FILE
- C. SKIP\_FILE\_ < num >
- D. ABORT\_STATEMENT

**Answer: B** ([LEAVE A REPLY](#))

The correct answer is B. SKIP\_FILE .

Important correction:

The provided answer A. CONTINUE is not correct for Snowpipe.

For regular bulk loading with COPY INTO < table > , the default ON\_ERROR behavior is commonly associated with ABORT\_STATEMENT. However, Snowpipe has different behavior. In Snowpipe, the default ON\_ERROR option is SKIP\_FILE .

Why B is correct:

With ON\_ERROR = SKIP\_FILE, if Snowpipe encounters an error while loading a file, the file is skipped instead of causing the entire continuous loading process to fail.

Why the other options are incorrect:

A). CONTINUE continues loading and skips individual error rows, but it is not the default Snowpipe behavior.

C). SKIP\_FILE\_ < num > skips a file after a specified number of errors, but it is not the default.

D). ABORT\_STATEMENT is associated with standard COPY INTO < table > behavior, not the default Snowpipe behavior.

Official Snowflake documentation reference:

Snowflake documentation for Snowpipe explains that the default ON\_ERROR copy option for Snowpipe is SKIP\_FILE.

Reference: Snowflake Documentation - Snowpipe; Snowflake Documentation - CREATE PIPE; Snowflake Documentation - COPY INTO < table > copy options; SnowPro Core Study Guide - Data Loading and Unloading.

==

### NEW QUESTION: 413

Which MINIMUM set of privileges is required to temporarily bypass an active network policy by configuring the user object property MINS\_TO\_BYPASS\_NETWORK\_POLICY?

A. Only while in the ACCOUNTADMIN role

B. Only while in the securityadmin role

C. Only the role with the ownership privilege on the network policy

D. Only Snowflake Support can set the value for this object property

**Answer:** [\(SHOW ANSWER\)](#)

To temporarily bypass an active network policy by configuring the user object property MINS\_TO\_BYPASS\_NETWORK\_POLICY, the minimum set of privileges required is having the ACCOUNTADMIN role. This role has the necessary privileges to make such changes, including modifying user properties that affect network policies.

Snowflake Documentation: Network Policy Management

### NEW QUESTION: 414

How often are encryption keys automatically rotated by Snowflake?

A. 30 Days

B. 60 Days

C. 90 Days

D. 365 Days

**Answer:** A [\(LEAVE A REPLY\)](#)

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: 415

What types of data listings are available in the Snowflake Data Marketplace? (Choose two.)

- A. Reader
- B. Consumer
- C. Vendor
- D. Standard
- E. Personalized

**Answer:** (SHOW ANSWER)

In the Snowflake Data Marketplace, the types of data listings available include 'Vendor', which refers to the providers of data, and 'Personalized', which indicates customized data offerings tailored to specific consumer needs<sup>45</sup>.

#### NEW QUESTION: 416

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

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

**Answer:** A,E (LEAVE A REPLY)

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

Role: ACCOUNTADMIN:

The ACCOUNTADMIN role has full administrative rights across the Snowflake account.

This role can manage all aspects of the Snowflake environment, including network policies.

Role with Global ATTACH POLICY Privilege:

Any role that has been granted the global ATTACH POLICY privilege can activate network policies.

This privilege allows the role to attach policies that control network access to the account.

Snowflake Documentation: Network Policies

#### NEW QUESTION: 417

How can a user optimize the performance of a COPY INTO < table > query that is being run on a very large file?

- A. Split the file into smaller files.
- B. Increase the virtual warehouse size.
- C. Set the SIZE\_LIMIT option to be larger than the file size.
- D. Set the COMPRESSION = AUTO option in the file format.

**Answer:** A (LEAVE A REPLY)

The correct answer is A. Split the file into smaller files .

Snowflake can load files in parallel. A single very large file limits parallelism because fewer worker threads can process the load concurrently. Splitting a very large file into multiple smaller files allows Snowflake to load the files in parallel, improving performance.

Why A is correct:

For best loading performance, Snowflake recommends splitting large data files into smaller files of a recommended size range so that the load operation can be parallelized.

Why the other options are incorrect:

B). Increasing the warehouse size can sometimes help, but if the load is constrained by a single very large file, file splitting is the better optimization.

C). SIZE\_LIMIT controls the maximum amount of data loaded during a COPY command. It does not improve performance for a very large file.

D). COMPRESSION = AUTO lets Snowflake detect compression type. It does not solve the lack of parallelism caused by a single large file.

Official Snowflake documentation reference:

Snowflake documentation recommends splitting large files into smaller files to increase parallelism and improve load performance.

Reference: Snowflake Documentation - Data loading considerations; Snowflake Documentation - COPY INTO < table > ; SnowPro Core Study Guide - Data Loading and Unloading.

#### **NEW QUESTION: 418**

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** ([LEAVE A REPLY](#))

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.

References = [ COF-C03 ] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

#### **NEW QUESTION: 419**

Which data types are supported by the Search Optimization Service? Select TWO.

**A.** INTEGER

**B.** GEOMETRY

**C.** FLOAT

**D.** VARIANT

**E.** BOOLEAN

**Answer: (SHOW ANSWER)**

The correct answers are A. INTEGER and B. GEOMETRY .

The Snowflake Search Optimization Service can improve performance for selective queries by maintaining a search access path. It supports several data types depending on the type of search being optimized, including scalar equality searches and geospatial searches.

Why A is correct:

INTEGER is a numeric data type. Search optimization supports equality searches on supported scalar data types, including numeric types such as integers.

Why B is correct:

GEOMETRY is supported for geospatial search optimization. Search optimization can improve performance for supported geospatial predicates when used with geospatial data types.

Why the other options are less appropriate in this question:

C). FLOAT is not the best answer in this option set. Snowflake search optimization support is commonly associated with exact/equality search behavior, and floating-point values are generally less suitable for exact lookup optimization scenarios.

D). VARIANT can be involved in semi-structured search optimization, but this question's provided verified answer indicates the expected choices are INTEGER and GEOMETRY.

E). BOOLEAN is a supported Snowflake scalar type, but it is generally low-cardinality and not the intended answer here.

Official Snowflake documentation reference:

Snowflake documentation explains that the Search Optimization Service supports specific query types and data types, including equality searches on supported scalar values and geospatial search optimization for geospatial data types such as GEOMETRY.

Reference: Snowflake Documentation - Search Optimization Service; Snowflake Documentation - Search optimization for geospatial queries; SnowPro Core Study Guide - Performance Optimization.

#### **NEW QUESTION: 420**

What is a best practice after creating a custom role?

**A.** Create the custom role using the SYSADMIN role.

**B.** Assign the custom role to the SYSADMIN role

**C.** Assign the custom role to the PUBLIC role

**D.** Add\_\_CUSTOM to all custom role names

**Answer: B (LEAVE A REPLY)**

Assigning the custom role to the SYSADMIN role is considered a best practice because it allows the SYSADMIN role to manage objects created by the custom role. This is important for maintaining proper access control and ensuring that the SYSADMIN can perform necessary administrative tasks on objects created by users with the custom role.

References:

[ COF-C03 ] SnowPro Core Certification Exam Study Guide

Section 1.3 - SnowPro Core Certification Study Guide1

#### **NEW QUESTION: 421**

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

**A.** True

**B.** False

**Answer: B (LEAVE A REPLY)**

Snowpipe via REST API can reference both named internal stages within Snowflake and external stages, such as Amazon S3, Google Cloud Storage, or Microsoft Azure1. This means that Snowpipe is not limited to only external stages as a source for data loading.

References = [ COF-C03 ] SnowPro Core Certification Exam Study Guide, Snowflake Documentation1 Reference:  
<https://community.snowflake.com/s/article/Making-Transient-table-by-Default>

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 422

What does an integration between Snowflake and Microsoft Private Link or AWS PrivateLink support?

- A. The isolation of data within a Snowflake account.
- B. A secure, direct connection to Snowflake that does not use the internet.
- C. The use of Secure Data Sharing among Snowflake accounts.
- D. A Virtual Private Network (VPN) between a user and Snowflake.

**Answer: B (LEAVE A REPLY)**

#### NEW QUESTION: 423

Users are responsible for data storage costs until what occurs?

- A. Data expires from Time Travel
- B. Data expires from Fail-safe
- C. Data is deleted from a table
- D. Data is truncated from a table

**Answer: B (LEAVE A REPLY)**

Users are responsible for data storage costs in Snowflake until the data expires from the Fail-safe period. Fail-safe is the final stage in the data lifecycle, following Time Travel, and provides additional protection against accidental data loss. Once data exits the Fail-safe state, users are no longer billed for its storage

#### NEW QUESTION: 424

Which table is designed to store a catalog of stage files in cloud storage?

- A. Iceberg table
- B. Directory table
- C. Dynamic table
- D. Hybrid table

**Answer: B (LEAVE A REPLY)**

#### NEW QUESTION: 425

Which type of URL gives permanent access to files in cloud storage?

- A. Account URL
- B. Pre-signed URL

C. Scoped URL

D. File URL

**Answer:** ([SHOW ANSWER](#))

#### **NEW QUESTION: 426**

Will data cached in a warehouse be lost when the warehouse is resized?

A. Possibly, if the warehouse is resized to a smaller size and the cache no longer fits.

B. Yes. because the compute resource is replaced in its entirety with a new compute resource.

C. No. because the size of the cache is independent from the warehouse size

D. Yes. because the new compute resource will no longer have access to the cache encryption key

**Answer:** C ([LEAVE A REPLY](#))

When a Snowflake virtual warehouse is resized, the data cached in the warehouse is not lost. This is because the cache is maintained independently of the warehouse size. Resizing a warehouse, whether scaling up or down, does not affect the cached data, ensuring that query performance is not impacted by such changes.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Virtual Warehouse Performance<sup>1</sup>

#### **NEW QUESTION: 427**

Which VALIDATION\_MODE value will return the errors across the files specified in a COPY command, including files that were partially loaded during an earlier load?

A. RETURN\_-1\_ROWS

B. RETURN\_n\_ROWS

C. RETURN\_ERRORS

D. RETURN ALL ERRORS

**Answer:** C ([LEAVE A REPLY](#))

The RETURN\_ERRORS value in the VALIDATION\_MODE option of the COPY command instructs Snowflake to validate the data files and return errors encountered across all specified files, including those that were partially loaded during an earlier load<sup>2</sup>. References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

#### **NEW QUESTION: 428**

What does a table with a clustering depth of 1 mean in Snowflake?

A. The table has only 1 micro-partition.

B. The table has 1 overlapping micro-partition.

C. The table has no overlapping micro-partitions.

D. The table has no micro-partitions.

**Answer:** ([SHOW ANSWER](#))

In Snowflake, a table 's clustering depth indicates the degree of micro-partition overlap based on the clustering keys defined for the table. A clustering depth of 1 implies that the table has no overlapping micro-partitions. This is an optimal scenario, indicating that the table 's data is well-clustered according to the specified clustering keys. Well-clustered data can lead to

more efficient query performance, as it reduces the amount of data scanned during query execution and improves the effectiveness of data pruning.

References:

Snowflake Documentation on Clustering: Understanding Clustering Depth

#### **NEW QUESTION: 429**

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

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

**Answer: D** ([LEAVE A REPLY](#))

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

#### **NEW QUESTION: 430**

Which type of workload is recommended for Snowpark-optimized virtual warehouses?

- A.** Workloads with ad hoc analytics
- B.** Workloads that have large memory requirements
- C.** Workloads with unpredictable data volumes for each query
- D.** Workloads that are queried with small table scans and selective filters

**Answer: B** ([LEAVE A REPLY](#))

Snowpark-optimized virtual warehouses in Snowflake are designed to efficiently handle workloads with large memory requirements. Snowpark is a developer framework that allows users to write code in languages like Scala, Java, and Python to process data in Snowflake. Given the nature of these programming languages and the types of data processing tasks they are typically used for, having a virtual warehouse that can efficiently manage large memory-intensive operations is crucial.

Understanding Snowpark-Optimized Virtual Warehouses:

Snowpark allows developers to build complex data pipelines and applications within Snowflake using familiar programming languages.

These virtual warehouses are optimized to handle the execution of Snowpark workloads, which often involve large datasets and memory-intensive operations.

Large Memory Requirements:

Workloads with large memory requirements include data transformations, machine learning model training, and advanced analytics.

These operations often need to process significant amounts of data in memory to perform efficiently.

Snowpark-optimized virtual warehouses are configured to provide the necessary memory resources to support these tasks, ensuring optimal performance and scalability.

Other Considerations:

While Snowpark can handle other types of workloads, its optimization for large memory tasks makes it particularly suitable for scenarios where data processing needs to be done in-memory.

Snowflake's ability to scale compute resources dynamically also plays a role in efficiently managing large memory workloads, ensuring that performance is maintained even as data volumes grow.

Snowflake Documentation: Introduction to Snowpark

Snowflake Documentation: Virtual Warehouses

### NEW QUESTION: 431

Which command will provide an account with access to a share?

- A. ALTER SHARE
- B. CREATE SHARE
- C. DESCRIBE SHARE
- D. GRANT < privilege > ... TO SHARE

**Answer: A ([LEAVE A REPLY](#))**

The correct answer is A. ALTER SHARE .

To provide a consumer account with access to an existing share, the provider adds the consumer account to the share. This is done using the ALTER SHARE command with the ADD ACCOUNTS clause.

Example:

```
ALTER SHARE my_share
```

```
ADD ACCOUNTS = consumer_org.consumer_account;
```

Why A is correct:

ALTER SHARE modifies an existing share. Adding consumer accounts to a share gives those accounts access to the share.

Why the other options are incorrect:

B). CREATE SHARE creates a new share object, but it does not by itself add a consumer account to an existing share.

C). DESCRIBE SHARE displays metadata about a share; it does not grant access.

D). GRANT < privilege > ... TO SHARE grants database object privileges to the share, such as usage on a database or select on a table. It does not add a consumer account to the share.

Official Snowflake documentation reference:

Snowflake documentation explains that ALTER SHARE ... ADD ACCOUNTS is used to add consumer accounts to a share.

Object privileges are granted separately to the share using GRANT ... TO SHARE.

Reference: Snowflake Documentation - ALTER SHARE; Snowflake Documentation - Secure Data Sharing; Snowflake Documentation - GRANT ... TO SHARE; SnowPro Core Study Guide - Data Protection and Governance.

==

### NEW QUESTION: 432

What is the purpose of key-pair authentication?

- A. To encrypt data that can be accessed with one of two keys
- B. To allow multiple users to securely share data
- C. To securely authenticate a user or an application
- D. To enforce Multi-Factor Authentication, or MFA

**Answer: ([SHOW ANSWER](#))**

The correct answer is C. To securely authenticate a user or an application .

Key-pair authentication uses a public/private key pair to authenticate to Snowflake. It is commonly used for service accounts, automation, applications, and client tools that need secure authentication without relying on passwords.

Why C is correct:

In key-pair authentication, the private key is used by the client, and the public key is associated with the Snowflake user. Snowflake verifies the authentication request using the public key.

Why the other options are incorrect:

A). Key-pair authentication is about authentication, not encrypting Snowflake data for access by one of two keys.

B). Secure data sharing is a separate Snowflake feature and is not the purpose of key-pair authentication.

D). MFA is a separate authentication control. Key-pair authentication does not itself enforce MFA.

Official Snowflake documentation reference:

Snowflake documentation describes key-pair authentication as an authentication method using a public /private key pair for users and applications.

Reference: Snowflake Documentation - Key-pair authentication; Snowflake Documentation - User authentication; SnowPro Core Study Guide - Security and Access Control.

==

### NEW QUESTION: 433

Which Snowflake function will interpret an input string as a JSON document, and produce a VARIANT value?

A. parse\_json()

B. json\_extract\_path\_text()

C. object\_construct()

D. flatten

**Answer: A (LEAVE A REPLY)**

The parse\_json() function in Snowflake interprets an input string as a JSON document and produces a VARIANT value containing the JSON document. This function is specifically designed for parsing strings that contain valid JSON information<sup>3</sup>.

### NEW QUESTION: 434

Which function is used to unload a relational table into a JSON file\*?

A. PARSE\_JSON

B. JSON\_EXTRACT\_PATH\_TEXT

C. OBJECT\_CONSTRUCT

D. TO\_JSON

**Answer: D (LEAVE A REPLY)**

The TO\_JSON function in Snowflake is used to convert a relational table or individual rows into JSON format.

This function is helpful for exporting data in JSON format.

Using TO\_JSON Function:

```
SELECT TO_JSON(OBJECT_CONSTRUCT(*))
```

```
FROM my_table;
```

Exporting Data: The TO\_JSON function converts the table rows into JSON format, which can then be exported to a file.

References:

Snowflake Documentation: TO\_JSON Function

**NEW QUESTION: 435**

Which task is supported by the use of Access History in Snowflake?

- A. Data backups
- B. Cost monitoring
- C. Compliance auditing
- D. Performance optimization

**Answer: C** ([LEAVE A REPLY](#))

Access History in Snowflake is primarily utilized for compliance auditing. The Access History feature provides detailed logs that track data access and modifications, including queries that read from or write to database objects. This information is crucial for organizations to meet regulatory requirements and to perform audits related to data access and usage.

Role of Access History: Access History logs are designed to help organizations understand who accessed what data and when. This is particularly important for compliance with various regulations that require detailed auditing capabilities.

How Access History Supports Compliance Auditing:

By providing a detailed log of access events, organizations can trace data access patterns, identify unauthorized access, and ensure that data handling complies with relevant data protection laws and regulations.

Access History can be queried to extract specific events, users, time frames, and accessed objects, making it an invaluable tool for compliance officers and auditors.

Reference: For more information on how Access History supports compliance auditing, refer to the Snowflake documentation on Access History: [https://docs.snowflake.com/en/sql-reference/account-usage/access\\_history.html](https://docs.snowflake.com/en/sql-reference/account-usage/access_history.html)

**NEW QUESTION: 436**

A user wants to access 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** ([LEAVE A REPLY](#))

A pre-signed URL provides a way to access objects stored in a stage without requiring authentication to Snowflake. This URL contains all the necessary information for access control, including an expiration time, and it's signed by the service's credentials. This mechanism allows users or applications to access specific files stored in a cloud storage location (stage) directly, bypassing the need to authenticate into Snowflake, ensuring secure and temporary access to the data. References: Snowflake Documentation on Stages and Data Loading

**corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

**NEW QUESTION: 437**

What virtual warehouse feature or setting will reduce the performance impact when running larger-than-average queries by offloading portions of the query processing work to shared compute resources?

- A. Using the search optimization service
- B. Using a multi-cluster virtual warehouse
- C. Using the query acceleration service
- D. Disabling auto-suspend

**Answer:** (SHOW ANSWER)

**NEW QUESTION: 438**

How should clustering be used to optimize the performance of queries that run on a very large table?

- A. Manually re-cluster the table regularly.
- B. Choose one high cardinality column as the clustering key.
- C. Use the column that is most-frequently used in query select clauses as the clustering key.
- D. Assess the average table depth to identify how clustering is impacting the query.

**Answer: B** (LEAVE A REPLY)

For optimizing the performance of queries that run on a very large table, it is recommended to choose one high cardinality column as the clustering key. This helps to co-locate similar rows in the same micro-partitions, improving scan efficiency in queries by skipping data that does not match filtering predicates.

References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 439**

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

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

**Answer: D** (LEAVE A REPLY)

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

**NEW QUESTION: 440**

A Snowflake Practitioner manages a database called SALES\_DB. This database has a PUBLIC schema and an ORDERS table that is updated frequently. The database has a data retention period of 7 days. On November 16, 2025, at 3:00 PM, the Practitioner accidentally ran a DELETE command that removed all records from the ORDERS table. Which statement will recover all deleted records?

- A. CREATE TABLE orders\_restored CLONE orders AT (TIMESTAMP = > ' 2025-11-16 10:00:00 ' :: TIMESTAMP);

- B. CREATE TABLE orders\_restored CLONE orders BEFORE (TIMESTAMP = > ' 2025-11-16 15:00:00' ::TIMESTAMP);
- C. CREATE TABLE orders\_restored CLONE orders BEFORE (OFFSET = > -60\*5);
- D. CREATE TABLE orders\_restored CLONE orders AT (TIMESTAMP = > ' 2025-11-16 15:00:00 ' ::TIMESTAMP);

**Answer:** ([SHOW ANSWER](#))

To recover data deleted at 3:00 PM, the clone should be created from a point immediately before the delete statement. BEFORE (TIMESTAMP = > ' 2025-11-16 15:00:00 ' ) returns the table state before that timestamp.

Why B is correct:

Snowflake Time Travel supports cloning a table at or before a specific timestamp within the retention period.

Reference: Snowflake Documentation - Time Travel; CREATE TABLE ... CLONE; AT | BEFORE.

==

#### NEW QUESTION: 441

Which operation can be performed on Snowflake external tables?

- A. INSERT
- B. JOIN
- C. RENAME
- D. ALTER

**Answer: B** ([LEAVE A REPLY](#))

Snowflake external tables are read-only, which means data manipulation language (DML) operations like INSERT, RENAME, or ALTER cannot be performed on them. However, external tables can be used for query and join operations<sup>3</sup>.

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

#### NEW QUESTION: 442

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\_ses s i on view

**Answer: B** ([LEAVE A REPLY](#))

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 Schema<sup>1</sup>

#### NEW QUESTION: 443

When cloning a schema, which Snowflake object will not be included in the clone?

- A. A named internal stage
- B. A User-Defined Function (UDF)

C. An external stage

D. A task

**Answer: C** ([LEAVE A REPLY](#))

#### NEW QUESTION: 444

Which functions are used to aggregate values into an array? Select TWO.

A. ARRAY\_AGG

B. ARRAY\_SLICE

C. ARRAY\_APPEND

D. ARRAY\_CONSTRUCT

E. ARRAY\_TO\_STRING

**Answer: A** ([LEAVE A REPLY](#))

The clearly correct answer from the listed options is A. ARRAY\_AGG .

Important correction:

The provided answer A, D is not fully correct if the question is asking about aggregate functions .

ARRAY\_CONSTRUCT creates an array from explicitly supplied values, but it does not aggregate values across multiple rows.

Why A is correct:

ARRAY\_AGG is an aggregate function. It aggregates input values from multiple rows into a single array.

Example:

```
SELECT ARRAY_AGG(product_id)
```

```
FROM orders;
```

This returns an array containing product\_id values from multiple rows.

Why D is not an aggregate function:

ARRAY\_CONSTRUCT creates an array from the values passed into the function.

Example:

```
SELECT ARRAY_CONSTRUCT( ' A ' , ' B ' , ' C ' );
```

This creates an array from listed expressions, but it does not aggregate values across rows.

Why the other options are incorrect:

B). ARRAY\_SLICE returns a subset of an existing array.

C). ARRAY\_APPEND appends a value to an existing array.

E). ARRAY\_TO\_STRING converts an array to a string.

Official Snowflake documentation reference:

Snowflake documentation describes ARRAY\_AGG as an aggregate function that returns the input values, pivoted into an array.

ARRAY\_CONSTRUCT is documented as a scalar function that constructs an array from input expressions.

Reference: Snowflake Documentation - ARRAY\_AGG; Snowflake Documentation - ARRAY\_CONSTRUCT; SnowPro Core Study Guide - Working with Semi-Structured Data.

#### NEW QUESTION: 445

Which role must be used to create resource monitors?

A. SECURITYADMIN

B. ACCOUNTADMIN

C. SYSADMIN

D. ORGADMIN

**Answer: B** ([LEAVE A REPLY](#))

In Snowflake, the ACCOUNTADMIN role is required to create resource monitors. Resource monitors are used to manage and monitor the consumption of compute resources. The ACCOUNTADMIN role has the necessary privileges to create, configure, and manage resource monitors across the account.

Snowflake Documentation: Resource Monitors

#### **NEW QUESTION: 446**

Which Snowflake tasks will take advantage of underlying micro-partition metadata? Select TWO.

A. Query pruning

B. Queries using a result cache

C. Queries using a local disk cache

D. Queries using a remote disk cache

E. Operations using Data Manipulation Language, or DML

**Answer: (**[SHOW ANSWER](#)**)**

The correct answers are A. Query pruning and E. Operations using Data Manipulation Language, or DML .

Snowflake automatically gathers and maintains metadata about micro-partitions. This metadata includes information such as the range of values in columns, distinct value counts, and other statistics. Snowflake uses this metadata to optimize query execution and table operations.

Why A is correct:

Query pruning is one of the most important uses of micro-partition metadata. Snowflake can skip micro-partitions that do not contain values needed by a query predicate.

Example:

```
SELECT *
```

```
FROM sales
```

```
WHERE sale_date = '2026-01-01' ;
```

If Snowflake knows from micro-partition metadata that certain micro-partitions do not contain `sale_date = '2026-01-01'` , those micro-partitions can be skipped.

Why E is correct:

DML operations such as UPDATE, DELETE, and MERGE can also benefit from micro-partition metadata.

Snowflake can use metadata to identify which micro-partitions are affected by the operation instead of scanning unnecessary partitions.

Why the other options are incorrect:

B). The result cache returns previously computed query results. It avoids re-executing the query and therefore does not rely on micro-partition pruning in the same way.

C). Local disk cache stores data on warehouse compute resources after access, but it is not the main feature that uses micro-partition metadata.

D). "Remote disk cache" is not the standard Snowflake concept tested here. Snowflake has remote storage, warehouse cache, and result cache, but query pruning is the metadata-driven feature.

Official Snowflake documentation reference:

Snowflake documentation explains that micro-partition metadata is used for efficient query pruning and optimization. It also supports efficient table maintenance and DML operations by helping Snowflake identify relevant micro-partitions.

Reference: Snowflake Documentation - Micro-partitions and data clustering; Snowflake Documentation - Query pruning; SnowPro Core Study Guide - Snowflake Architecture.

==

#### NEW QUESTION: 447

Which table type exists only within the session it was created and is only visible to the user who created it?

- A. Hybrid
- B. Dynamic
- C. Transient
- D. Temporary

**Answer: (SHOW ANSWER)**

The correct answer is D. Temporary .

A temporary table exists only for the duration of the session in which it was created. It is only visible to the user session that created it and is automatically dropped when the session ends.

Why D is correct:

Temporary tables are useful for session-specific intermediate data. They do not persist beyond the session and are not visible to other users or sessions.

Example:

```
CREATE TEMPORARY TABLE temp_sales AS
SELECT *
FROM sales
WHERE sales_date = CURRENT_DATE;
```

Why the other options are incorrect:

- A). Hybrid tables are optimized for transactional workloads and are not session-only objects.
- B). Dynamic tables are used for declarative, automated data transformation and are persistent database objects.
- C). Transient tables persist until explicitly dropped and are visible according to granted privileges. They do not have Fail-safe, but they are not session-only.

Official Snowflake documentation reference:

Snowflake documentation describes temporary tables as existing only within the session in which they are created and being dropped automatically at the end of the session.

Reference: Snowflake Documentation - Temporary tables; Snowflake Documentation - Table types; SnowPro Core Study Guide - SQL and Snowflake Objects.

#### NEW QUESTION: 448

A JSON document is stored in the source\_column of type VARIANT. The document has an array called elements. The array contains the name key that has a string value How can a Snowflake user extract the name from the first element?

- A. Source\_column.element[1]:name
- B. Source\_column.element[0]:name
- C. Source\_column.element[1].name

**D.** Source\_column.element[0]:name

**Answer: C** ([LEAVE A REPLY](#))

In Snowflake, when dealing with semi-structured data such as a JSON document stored in a VARIANT column, the proper syntax to extract a value is to use the column name followed by the path to the specific element. Since arrays in JSON are zero-indexed, the first element is referenced with [0]. Therefore, to extract the name from the first element of the elements array, the correct syntax is Source\_column:elements[0].name.

References: Snowflake Documentation on Semi-Structured Data

#### **NEW QUESTION: 449**

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** ([LEAVE A REPLY](#))

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: 450**

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** ([LEAVE A REPLY](#))

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: 451**

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

**A.** GO driver

**B.** Node.js driver

**C.** ODBC driver

**D.** Python connector

**E.** Spark connector

**Answer:** ([SHOW ANSWER](#))

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

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 452

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 (LEAVE A REPLY)**

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: 453

A virtual warehouse query queuing time has been increasing due to a large number of users accessing the warehouse at the same time each day. Which warehouse configurations would address this issue MOST cost-effectively? (Select TWO).

- A. Run the warehouse in Auto-scale mode.
- B. Scale-up the warehouse.
- C. Run the warehouse using an Economy policy.
- D. Run the warehouse in Maximized mode.
- E. Scale-out the warehouse.

**Answer: A,E (LEAVE A REPLY)**

#### NEW QUESTION: 454

What is a fundamental characteristic of Snowflake micro-partitions?

- A. They are sized based on Time Travel requirements.
- B. They can be read directly as files.

- C. Once established they cannot be changed.
- D. They serve as an index for Snowflake tables.

**Answer: A** ([LEAVE A REPLY](#))

#### NEW QUESTION: 455

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** ([LEAVE A REPLY](#))

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:

Current Schema: Snowflake first checks if an object with the given name exists in the schema currently in use for the session.

PUBLIC Schema: If the object isn't found in the current schema, Snowflake looks in the PUBLIC schema of the current database.

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: 456

Which of the following features are available with the Snowflake Enterprise edition? (Choose two.)

- A. Database replication and failover
- B. Automated index management
- C. Customer managed keys (Tri-secret secure)
- D. Extended time travel
- E. Native support for geospatial data

**Answer: A,D** ([LEAVE A REPLY](#))

The Snowflake Enterprise edition includes database replication and failover for business continuity and disaster recovery, as well as extended time travel capabilities for longer data retention periods<sup>1</sup>.

#### NEW QUESTION: 457

Which privilege must be granted to a share to allow secure views the ability to reference data in multiple databases?

- A. CREATE\_SHARE on the account
- B. SHARE on databases and schemas
- C. SELECT on tables used by the secure view
- D. REFERENCE\_USAGE on databases

**Answer: D** ([LEAVE A REPLY](#))

To allow secure views the ability to reference data in multiple databases, the `REFERENCE_USAGE` privilege must be granted on each database that contains objects referenced by the secure view<sup>2</sup>. This privilege is necessary before granting the `SELECT` privilege on a secure view to a share.

#### **NEW QUESTION: 458**

Which Snowflake governance feature allows users to assign metadata labels to improve data governance and database access control?

- A.** Secure functions
- B.** Secure views
- C.** Object tagging
- D.** Row-level security

**Answer:** [\(SHOW ANSWER\)](#)

Object tagging in Snowflake allows users to assign metadata labels to various database objects to improve data governance and access control. Tags can be used to categorize and manage data based on business needs, helping to enforce governance policies and streamline database administration.

References:

Snowflake Documentation: Object Tagging

#### **NEW QUESTION: 459**

What are the benefits of using the Snowflake `Access_History` view? (Select TWO)

- A.** It allows users to track the relationships among objects.
- B.** It identifies data that is not being used so that it can be archived or deleted.
- C.** It tracks how sensitive data has been moved.
- D.** It audits whether query policies are being enforced correctly.
- E.** It allows the user to compare query performance data over time.

**Answer:** **B,C** [\(LEAVE A REPLY\)](#)

#### **NEW QUESTION: 460**

How does the Snowflake search optimization service improve query performance?

- A.** It improves the performance of range searches.
- B.** It defines different clustering keys on the same source table.
- C.** It improves the performance of all queries running against a given table.
- D.** It improves the performance of equality searches.

**Answer:** **D** [\(LEAVE A REPLY\)](#)

The Snowflake Search Optimization Service is designed to enhance the performance of specific types of queries on large tables. The correct answer is:

**D** . It improves the performance of equality searches: The service optimizes the performance of queries that use equality search conditions (e.g., `WHERE column = value`). It creates and maintains a search index on the table's columns, which significantly speeds up the retrieval of rows based on those equality search conditions.

This optimization is particularly beneficial for large tables where traditional scans might be inefficient for equality searches. By using the Search Optimization Service, Snowflake can leverage the search indexes to quickly locate the rows that match the search criteria without scanning the entire table.

References:

Snowflake Documentation: Search Optimization Service at [Snowflake Documentation](#)

#### NEW QUESTION: 461

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

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

**Answer: A,C (LEAVE A REPLY)**

Snowflake network policies are a security feature that allows administrators to control access to Snowflake by specifying allowed and blocked IP address ranges. These policies apply to all editions of Snowflake, making them widely applicable across different Snowflake environments. They are specifically designed to restrict or enable access based on the originating IP addresses of client requests, adding an extra layer of security.

Network policies are not applied to roles but are set at the account or user level. They are not activated using ALTER DATABASES SQL commands but are managed through ALTER ACCOUNT or ALTER NETWORK POLICY commands. The management of network policies does not exclusively require the ORGADMIN role; instead, they can be managed by users with the necessary privileges on the account.

Reference: Snowflake Documentation on Network Policies: <https://docs.snowflake.com/en/user-guide/network-policies.html>

#### NEW QUESTION: 462

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

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

**Answer: (SHOW ANSWER)**

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

#### NEW QUESTION: 463

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

- A. Include the SQL statements in the body of the task CREATE TASK mytask .. AS INSERT INTO target1 SELECT .. FROM stream\_s1 WHERE .. INSERT INTO target2 SELECT .. FROM stream\_s1
- B. A stored procedure can have only one DML statement per stored procedure invocation and therefore the user should sequence stored procedure calls in the task definition CREATE TASK mytask .... AS

- C.** Use a stored procedure executing multiple SQL statements and invoke the stored procedure from the task. CREATE TASK mytask .... AS call stored\_proc\_multiple\_statements\_inside();
- D.** Create a task for each SQL statement (e.g. resulting in task1, task2, etc.) and string the series of SQL statements by having a control task calling task1, task2, etc. sequentially.

**Answer: C** ([LEAVE A REPLY](#))

To execute a series of SQL statements using a task, a user would use a stored procedure that contains multiple SQL statements and invoke this stored procedure from the task. References: Snowflake Documentation<sup>2</sup>.

#### NEW QUESTION: 464

What is the MINIMUM Snowflake edition that supports database replication?

- A.** Standard
- B.** Enterprise
- C.** Business Critical
- D.** Virtual Private Snowflake (VPS)

**Answer: B** ([LEAVE A REPLY](#))

The minimum Snowflake edition that supports database replication is the Enterprise edition. Database replication allows data to be replicated between different Snowflake accounts or regions, providing high availability and disaster recovery capabilities.

References:

Snowflake Documentation: Database Replication

#### NEW QUESTION: 465

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: A,C** ([LEAVE A REPLY](#))

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: 466

Which Snowflake tool or feature enables governed sharing of datasets with a selected group of consumers that are invited, internal or external, to a Snowflake user 's company?

- A.** Snowflake Marketplace
- B.** Snowpipe
- C.** Replication
- D.** Data Exchange

**Answer:** ([SHOW ANSWER](#))

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

**NEW QUESTION: 467**

Use of which feature or setting consumes Snowflake credits, but generates no storage costs?

- A. Table clustering
- B. Materialized view
- C. Search optimization service
- D. Query acceleration service

**Answer:** ([SHOW ANSWER](#))

**NEW QUESTION: 468**

How long can a data consumer who has a pre-signed URL access data files using Snowflake?

- A. Indefinitely
- B. Until the result\_cache expires
- C. Until the retention\_time is met
- D. Until the expiration time is exceeded

**Answer: C** ([LEAVE A REPLY](#))

A data consumer who has a pre-signed URL can access data files using Snowflake until the expiration time is exceeded. The expiration time is set when the pre-signed URL is generated and determines how long the URL remains valid.

**NEW QUESTION: 469**

What is the primary purpose of using a masking policy in Snowflake?

- A. To protect sensitive data from unauthorized access when queries are run.
- B. To automatically encrypt sensitive data when data is stored in Snowflake.
- C. To protect multiple columns that have different data types in a given table.
- D. To protect both column-level and row-level data.

**Answer: A** ([LEAVE A REPLY](#))

Masking policies in Snowflake are designed to protect sensitive information by dynamically hiding or obfuscating data based on the role of the user executing the query. This helps enforce data privacy and security by allowing only authorized users to see sensitive information. Masking policies do not encrypt data but apply rules to limit data visibility, ensuring sensitive data is protected during query execution without altering the underlying data.

**NEW QUESTION: 470**

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

- A. Network policy
- B. End-to-end encryption

- C. Multi-Factor Authentication (MFA)
- D. Role-Based Access Control (RBAC)

**Answer: A** ([LEAVE A REPLY](#))

A network policy in Snowflake is used to restrict access to the Snowflake account from unauthorized or unknown sources. It allows administrators to specify allowed IP address ranges, thus preventing access from any IP addresses not listed in the policy<sup>1</sup>.

#### NEW QUESTION: 471

Which Snowflake tool would be BEST to troubleshoot network connectivity?

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

**Answer: D** ([LEAVE A REPLY](#))

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

#### NEW QUESTION: 472

Which stage type can be altered and dropped?

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

**Answer: B** ([LEAVE A REPLY](#))

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: 473

A Snowflake user is investigating a slow-running query using the Query Profile. What would be a reason to increase the size of the virtual warehouse?

- A. Percentage Scanned from Cache: 81.22%
- B. Processing: 51.8%
- C. Remote Disk I/O: 40.6%
- D. Bytes Scanned: 3.24 GB

**Answer: B** ([LEAVE A REPLY](#))

The correct answer is B. Processing: 51.8% .

A high percentage of query time spent in Processing indicates that the query is spending significant time on CPU-intensive work, such as joins, aggregations, sorting, window functions, or expression evaluation.

Increasing the virtual warehouse size can provide more compute resources and may reduce processing time.

Why B is correct:

If the Query Profile shows a large portion of time spent in processing, scaling up the warehouse can help because a larger warehouse provides more compute power.

Why the other options are incorrect:

A). A high percentage scanned from cache is generally positive. It means much of the data was served from cache instead of remote storage.

C). High remote disk I/O may indicate data is being read from remote storage. Increasing warehouse size may not be the first or best fix; improving pruning, clustering, filters, or reducing scanned data may be more appropriate.

D). Bytes Scanned: 3.24 GB by itself is not enough to justify increasing warehouse size. The reason to scale up is usually compute pressure, not simply a raw scan number.

Official Snowflake documentation reference:

Snowflake documentation explains that Query Profile helps identify time spent in processing, local disk I/O, remote disk I/O, synchronization, and other execution categories. Increasing warehouse size can help when queries are constrained by compute processing.

Reference: Snowflake Documentation - Query Profile; Snowflake Documentation - Virtual warehouse performance; SnowPro Core Study Guide - Performance Optimization.

#### **NEW QUESTION: 474**

Which of the following describes the Snowflake Cloud Services layer?

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

**Answer:** ([SHOW ANSWER](#))

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

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

#### **NEW QUESTION: 475**

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:** ([SHOW ANSWER](#))

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: 476**

How is Single Sign-On (SSO) authentication used in Snowflake?

- A.** SSO is an integrated Snowflake feature powered by the Duo Security service which is managed completely by Snowflake.
- B.** SSO is an authentication method which uses a pair of keys, a public key and a private key. to verify the identity of a user.
- C.** SSO is an authentication method that uses a username and password in the API request header.
- D.** SSO is an authentication method that allows a user to sign into multiple applications with a single set of credentials.

**Answer:** ([SHOW ANSWER](#))

#### **NEW QUESTION: 477**

Which function will generate a URL that can be accessed by a non-Snowflake user?

- A.** BUILD\_SCOPED\_FILE\_URL
- B.** BUILD\_STAGE\_FILE\_URL
- C.** GET\_PREIGNED\_URL
- D.** GET\_STAGE\_LOCATION

**Answer:** **D** ([LEAVE A REPLY](#))

#### **NEW QUESTION: 478**

A Medium virtual warehouse processes scheduled reports for a team. More than 30 reports run simultaneously during business hours, and the team requires high availability, consistent performance, and minimal startup time.

What action should be taken to help avoid queuing when these reports are run?

- A.** Increase the size of the warehouse.
- B.** Enable a multi-cluster warehouse with Auto-scale mode and an Economy scaling policy.
- C.** Enable a multi-cluster warehouse with Maximized mode.
- D.** Enable a multi-cluster warehouse with Auto-scale mode and a Standard scaling policy.

**Answer:** **D** ([LEAVE A REPLY](#))

A multi-cluster warehouse in Auto-scale mode automatically starts additional clusters when concurrency increases and queries begin to queue. The Standard scaling policy prioritizes responsiveness and throughput by favoring the rapid startup of additional clusters, whereas the Economy policy prioritizes credit conservation and can allow queries to remain queued longer. This makes option D appropriate for a reporting workload that experiences a significant concurrency increase during business hours. Maximized mode keeps a fixed number of clusters running and is better suited to consistently high workloads that do not fluctuate significantly. Increasing the warehouse size scales up each cluster and can accelerate individual complex queries, but it does not address concurrency as directly as scaling out with additional clusters.

#### **NEW QUESTION: 479**

How does the search optimization service improve query performance?

- A.** By clustering the tables
- B.** By creating a persistent data structure
- C.** By using caching

**D.** By optimizing the use of micro-partitions

**Answer: B** ([LEAVE A REPLY](#))

The Search Optimization Service in Snowflake enhances query performance by creating a persistent data structure that enables faster access to specific data, particularly for queries with selective filters on columns not often used in clustering. This persistent structure accelerates data retrieval without depending on clustering or caching, thereby improving response times for targeted queries.

Snowflake's micro-partitioning automatically manages table structure, but search optimization allows further enhancement for certain high-frequency, specific access patterns.

#### **NEW QUESTION: 480**

What kind of authentication do Snowpipe REST endpoints use?

**A.** OAuth

**B.** Key-based

**C.** Username and password

**D.** Single Sign-On (SSO)

**Answer: (**[SHOW ANSWER](#)**)**

Snowpipe uses key-based authentication for its REST endpoints. This involves generating and using a key pair (public and private keys) to securely authenticate API requests.

Generate Key Pair: Generate a public and private key pair.

Register Public Key: Register the public key with the Snowflake user that will be making the API requests.

Authenticate Requests: Use the private key to sign API requests sent to Snowpipe REST endpoints.

References:

Snowflake Documentation: Key Pair Authentication & Key Rotation

Snowflake Documentation: Using Snowpipe REST API

#### **NEW QUESTION: 481**

A user runs a query on a large table using a Medium virtual warehouse. The same query is then run on the same table by a different user using a different warehouse.

Which caches will be used?

**A.** Local disk cache and metadata cache

**B.** Query result cache and warehouse cache

**C.** Warehouse cache and metadata cache

**D.** Metadata cache and query result cache

**Answer: D** ([LEAVE A REPLY](#))

The warehouse data cache is local to the virtual warehouse that populated it, so a different warehouse cannot reuse the first warehouse's cached table data. Persisted query results are independent of a specific warehouse and can be reused for an identical SELECT when the data is unchanged, the result is still available, reuse is enabled, and the executing role has the required privileges. Snowflake's Cloud Services layer also manages metadata and query caching outside individual warehouse-local storage.

Consequently, among the available options, metadata cache and query result cache are applicable, making D correct. A different user does not automatically prevent result reuse for a SELECT. The user's active role must have the necessary privileges on all objects referenced by the cached query.

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (1187 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

#### NEW QUESTION: 482

How does Snowflake ensure the integrity and consistency of data stored in tables?

- A. Dynamic Data Masking
- B. Row-level encryption
- C. Constraints
- D. Time Travel

**Answer: C** ([LEAVE A REPLY](#))

#### NEW QUESTION: 483

Which data formats are supported by Snowflake when unloading semi-structured data? (Select TWO).

- A. Binary file in Avro
- B. Binary file in Parquet
- C. Comma-separated JSON
- D. Newline Delimited JSON
- E. Plain text file containing XML elements

**Answer: (SHOW ANSWER)**

Snowflake supports a variety of file formats for unloading semi-structured data, among which Parquet and Newline Delimited JSON (NDJSON) are two widely used formats.

B). Binary file in Parquet: Parquet is a columnar storage file format optimized for large-scale data processing and analysis. It is especially suited for complex nested data structures.

D). Newline Delimited JSON (NDJSON): This format represents JSON records separated by newline characters, facilitating the storage and processing of multiple, separate JSON objects in a single file.

These formats are chosen for their efficiency and compatibility with data analytics tools and ecosystems, enabling seamless integration and processing of exported data.

References:

Snowflake Documentation: Data Unloading

#### NEW QUESTION: 484

Will data cached in a warehouse be lost when the warehouse is resized?

- A. Possibly, if the warehouse is resized to a smaller size and the cache no longer fits.

- B.** Yes. because the compute resource is replaced in its entirety with a new compute resource.
- C.** No. because the size of the cache is independent from the warehouse size
- D.** Yes. because the new compute resource will no longer have access to the cache encryption key

**Answer: C (LEAVE A REPLY)**

When a Snowflake virtual warehouse is resized, the data cached in the warehouse is not lost. This is because the cache is maintained independently of the warehouse size. Resizing a warehouse, whether scaling up or down, does not affect the cached data, ensuring that query performance is not impacted by such changes.

References:

[ COF-C03 ] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Virtual Warehouse Performance<sup>1</sup>

#### **NEW QUESTION: 485**

What is the recommended compressed file size range for continuous data loads using Snowpipe?

- A.** 8-16 MB
- B.** 16-24 MB
- C.** 10-99 MB
- D.** 100-250 MB

**Answer: D (LEAVE A REPLY)**

For continuous data loads using Snowpipe, the recommended compressed file size range is between 100-250 MB. This size range is suggested to optimize the number of parallel operations for a load and to avoid size limitations, ensuring efficient and cost-effective data loading.

#### **NEW QUESTION: 486**

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

- A.** Root task may be executed concurrently (i.e. multiple instances), it is recommended to leave some margins in the execution window to avoid missing instances of execution
- B.** Querying (select) the size of the stream content would help determine the warehouse size. For example, if querying large stream content, use a larger warehouse size
- C.** 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
- D.** 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 (LEAVE A REPLY)**

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-C03 ] SnowPro Core Certification Exam Study Guide

#### **NEW QUESTION: 487**

What is a machine learning and data science partner within the Snowflake Partner Ecosystem?

- A.** Informatica

- B. Power BI
- C. Adobe
- D. Data Robot

**Answer: D (LEAVE A REPLY)**

Data Robot is recognized as a machine learning and data science partner within the Snowflake Partner Ecosystem. It provides an enterprise AI platform that enables users to build and deploy accurate predictive models quickly. As a partner, Data Robot integrates with Snowflake to enhance data science capabilities<sup>2</sup>.

References:

[ COF-C03 ] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Machine Learning & Data Science Partners

<https://docs.snowflake.com/en/user-guide/ecosystem-analytics.html>

#### **NEW QUESTION: 488**

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

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

**Answer: C (LEAVE A REPLY)**

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

#### **NEW QUESTION: 489**

Which view or feature can be used to analyze query performance and identify which operations are the most time-consuming?

- A. Search Optimization Service
- B. QUERY\_HISTORY view
- C. Query Profile
- D. WAREHOUSE\_LOAD\_HISTORY view

**Answer: C (LEAVE A REPLY)**

Query Profile displays the execution plan for an individual query as a set of operator nodes. Each node represents an operation such as a table scan, filter, join, aggregate, sort, or data unload and shows the percentage of the query duration consumed by that operation. The Most Expensive Nodes pane identifies the operators that required the greatest execution time, allowing users to focus on the primary bottlenecks. QUERY\_HISTORY provides statement-level timing, warehouse, status, and resource metrics but does not by itself provide the same operator-level graphical analysis.

WAREHOUSE\_LOAD\_HISTORY helps evaluate warehouse load and queuing, while Search Optimization Service is a performance feature rather than a diagnostic view. Therefore, Query Profile is the appropriate tool for determining which individual operations consume the most execution time.

**Valid COF-C03 Dumps** shared by PrepPdf.com for Helping Passing COF-C03 Exam! PrepPdf.com now offer the **newest COF-C03 exam dumps**, the PrepPdf.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com COF-C03 dumps with Test Engine here: <https://www.preppdf.com/Snowflake/COF-C03-prepaway-exam-dumps.html> (**1187** Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)