



"Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns."

www.atmicnetworks.com

Warning: Keep connected with our support team for latest updates

Question: 1

The fail-safe retention period is how many days?

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

Answer: B

Explanation:

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

References: [Understanding and viewing Fail-safe | Snowflake Documentation](#)

Question: 2

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

- A. True
- B. False

Answer: A

Explanation:

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

Question: 3

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

Explanation:

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

Question: 4

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: B, C, D

Explanation:

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

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

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

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

[References: The information is verified according to the SnowPro Core Certification Study Guide and Snowflake documentation](#)

Question: 5

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

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

Answer: A

Explanation:

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

Question: 6

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

- A. True
- B. False

Answer: B

Explanation:

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

environment only.

Question: 7

True or False: Loading data into Snowflake requires that source data files be no larger than 16MB.

- A. True
- B. False

Answer: B

Explanation:

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](#)

Question: 8

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

- A. True
- B. False

Answer: A

Explanation:

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

Reference: [https://docs.snowflake.com/en/user-guide/warehouses-tasks.html#effects-of-resizing-a-suspended-](https://docs.snowflake.com/en/user-guide/warehouses-tasks.html#effects-of-resizing-a-suspended-warehouse)

[warehouse](#)

Question: 9

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

- A. True
- B. False

Answer: B

Explanation:

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

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

Question: 10

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

- A. Via the Snowflake Web Interface (UI)
- B. Via the data_share=true parameter
- C. Via SQL commands
- D. Via Virtual Warehouses

Answer: A, C

Explanation:

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

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

Via SQL commands: Snowflake also allows the creation and management of shares using SQL commands. This method is more suited for users who prefer scripting or need to automate the process.

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

Question: 11

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

A. True

B. False

Answer: B

Explanation:

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

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

Question: 12

True or False: It is possible for a user to run a query against the query result cache without requiring an active Warehouse.

A. True

B. False

Answer: A

Explanation:

Snowflake's architecture allows for the use of a query result cache that stores the results of queries for a period of time. If the same query is run again and the underlying data has not changed, Snowflake can retrieve the result from this cache without needing to re-run the query on an active warehouse, thus saving on compute resources.

Question: 13

A virtual warehouse's auto-suspend and auto-resume settings apply to which of the following?

- A. The primary cluster in the virtual warehouse
- B. The entire virtual warehouse
- C. The database in which the virtual warehouse resides
- D. The Queries currently being run on the virtual warehouse

Answer: B

Explanation:

The auto-suspend and auto-resume settings in Snowflake apply to the entire virtual warehouse. These settings allow the warehouse to automatically suspend when it's not in use, helping to save on compute costs. When queries or tasks are submitted to the warehouse, it can automatically resume operation. This functionality is designed to optimize resource usage and cost-efficiency.

References:

SnowPro Core Certification Exam Study Guide (as of 2021)

Snowflake documentation on virtual warehouses and their settings (as of 2021)

Question: 14

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: C, E

Explanation:

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

References:

[Continuous Data Protection | Snowflake Documentation1](#)

[Data Storage Considerations | Snowflake Documentation2](#)

[Snowflake SnowPro Core Certification Study Guide3](#)

[Snowflake Data Cloud Glossary](#)

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

Question: 15

Which of the following conditions must be met in order to return results from the results cache? (Select TWO).

- A. The user has the appropriate privileges on the objects associated with the query
- B. Micro-partitions have been reclustered since the query was last run
- C. The new query is run using the same virtual warehouse as the previous query
- D. The query includes a User Defined Function (UDF)
- E. The query has been run within 24 hours of the previously-run query

Answer: A, E

Explanation:

To return results from the results cache in Snowflake, certain conditions must be met:

Privileges: The user must have the appropriate privileges on the objects associated with the query.

This ensures that only authorized users can access cached data.

Time Frame: The query must have been run within 24 hours of the previously-run query. Snowflake's results cache is designed to store the results of queries for a short period, typically 24 hours, to improve performance for repeated queries.

Question: 16

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

Explanation:

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>

Question: 17

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

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

Answer: B

Explanation:

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

References:

[Snowflake Documentation on CREATE MATERIALIZED 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%20pl>

[ease%20contact%20Snowflake%20Support.](#)

Question: 18

What happens to the underlying table data when a CLUSTER BY clause is added to a Snowflake table?

- A. Data is hashed by the cluster key to facilitate fast searches for common data values
- B. Larger micro-partitions are created for common data values to reduce the number of partitions that must be scanned
- C. Smaller micro-partitions are created for common data values to allow for more parallelism
- D. Data may be colocated by the cluster key within the micro-partitions to improve pruning performance

Answer: D

Explanation:

When a CLUSTER BY clause is added to a Snowflake table, it specifies one or more columns to organize the data within the table's micro-partitions. This clustering aims to colocate data with similar values in the same or adjacent micro-partitions. By doing so, it enhances the efficiency of query pruning, where the Snowflake query optimizer can skip over irrelevant micro-partitions that do not contain the data relevant to the query, thereby improving performance.

References:

[Snowflake Documentation on Clustering Keys & Clustered Tables1.](#)

[Community discussions on how source data's ordering affects a table with a cluster key](#)

Question: 19

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

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

Answer: A

Explanation:

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

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

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

Question: 20

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

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

Answer: B, D, E

Explanation:

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

Question: 21

During periods of warehouse contention which parameter controls the maximum length of time a warehouse will hold a query for processing?

- A. STATEMENT_TIMEOUT IN SECONDS
- B. STATEMENT_QUEUED_TIMEOUT_IN_SECONDS
- C. MAX_CONCURRENCY_LEVEL
- D. QUERY_TIMEOUT_IN_SECONDS

Answer: B

Explanation:

The parameter STATEMENT_QUEUED_TIMEOUT_IN_SECONDS sets the limit for a query to wait in the queue in order to get its chance of running on the warehouse. The query will quit after reaching this limit. By default, the value of this parameter is 0 which means the queries will wait indefinitely in the waiting queue

https://community.snowflake.com/s/article/Warehouse-Concurrency-and-Statement-Timeout-Parameters#:~:text=The%20parameter%20STATEMENT_QUEUED_TIMEOUT_IN_SECONDS%20sets%20the,indefinitely%20in%20the%20waiting%20queue.

Question: 22

Which of the following indicates that it may be appropriate to use a clustering key for a table? (Select TWO).

- A. The table contains a column that has very low cardinality
- B. DML statements that are being issued against the table are blocked
- C. The table has a small number of micro-partitions
- D. Queries on the table are running slower than expected
- E. The clustering depth for the table is large

Answer: D, E

Explanation:

A clustering key in Snowflake is used to co-locate similar data within the same micro-partitions to improve query performance, especially for large tables where data is not naturally ordered or has become fragmented due to extensive DML operations. The appropriate use of a clustering key can lead to improved scan efficiency and better column compression, resulting in faster query execution times.

The indicators that it may be appropriate to use a clustering key for a table include:

D. Queries on the table are running slower than expected: This can happen when the data in the table is not well-clustered, leading to inefficient scans during query execution.

E. The clustering depth for the table is large: A large clustering depth indicates that the table's data is spread across

many micro-partitions, which can degrade query performance as more data needs to be scanned.

References:

[Snowflake Documentation on Clustering Keys & Clustered Tables](#)

[Snowflake Documentation on SYSTEM\\$CLUSTERING_INFORMATION](#)

[Stack Overflow discussion on cluster key selection in Snowflake](#)

Question: 23

Which Snowflake object enables loading data from files as soon as they are available in a cloud storage location?

- A. Pipe
- B. External stage
- C. Task
- D. Stream

Answer: A

Explanation:

In Snowflake, a Pipe is the object designed to enable the continuous, near-real-time loading of data from files as soon as they are available in a cloud storage location. Pipes use Snowflake's COPY command to load data and can be associated with a Stage object to monitor for new files. When new data files appear in the stage, the pipe automatically loads the data into the target table.

References:

[Snowflake Documentation on Pipes](#)

[SnowPro® Core Certification Study Guide](#)

<https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro.html>

Question: 24

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.MYSHEMA TO USER USER1;
```

C.

```
GRANT ROLE MYROLE TO USER USER1;
```

```
CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSHEMA TO USER1;
```

D.

```
GRANT ROLE MYROLE TO USER USER1;
```

```
CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSHEMA TO MYROLE;
```

Answer: D

Explanation:

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

The statement `GRANT ROLE MYROLE TO USER USER1;` grants the specified role to the user, allowing them to assume that role and exercise its privileges. The subsequent statement `CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSHEMA 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.MYSHEMA`.

References:

[Snowflake Documentation on Roles](#)

[Snowflake Documentation on Materialized Views](#)

Question: 25

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

A. UTF-8

B. UTF-16

C. ISO 859-1

D. ANSI_X3.A

Answer: A

Explanation:

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

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

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

Question: 26

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 cache

Answer: D

Explanation:

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](#)

Question: 27

Which cache type is used to cache data output from SQL queries?

- A. Metadata cache
- B. Result cache
- C. Remote cache
- D. Local file cache

Answer: B

Explanation:

The Result cache is used in Snowflake to cache the data output from SQL queries. This feature is designed to improve performance by storing the results of queries for a period of time. When the same or similar query is executed again, Snowflake can retrieve the result from this cache instead of re-computing the result, which saves time and computational resources.

References:

[Snowflake Documentation on Query Results Cache](#)

[SnowPro® Core Certification Study Guide](#)

Question: 28

What is a key feature of Snowflake architecture?

- A. Zero-copy cloning creates a mirror copy of a database that updates with the original
- B. Software updates are automatically applied on a quarterly basis
- C. Snowflake eliminates resource contention with its virtual warehouse implementation
- D. Multi-cluster warehouses allow users to run a query that spans across multiple clusters
- E. Snowflake automatically sorts DATE columns during ingest for fast retrieval by date

Answer: C

Explanation:

One of the key features of Snowflake's architecture is its unique approach to eliminating resource contention through the use of virtual warehouses. This is achieved by separating storage and compute resources, allowing multiple virtual warehouses to operate independently on the same data without affecting each other. This means that different workloads, such as loading data, running queries, or performing complex analytics, can be processed simultaneously without any performance degradation due to resource contention.

References:

[Snowflake Documentation on Virtual Warehouses](#)

[SnowPro® Core Certification Study Guide](#)

Question: 29

What is a limitation of a Materialized View?

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

Answer: D

Explanation:

Materialized Views in Snowflake are designed to store the result of a query and can be refreshed to maintain up-to-date data. However, they have certain limitations, one of which is that they cannot be defined using a JOIN clause.

This means that a Materialized View can only be created based on a single source table and cannot combine data from multiple tables using JOIN operations.

References:

[Snowflake Documentation on Materialized Views](#)

[SnowPro® Core Certification Study Guide](#)

Question: 30

What features does Snowflake Time Travel enable?

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

Answer: B, C

Explanation:

Snowflake Time Travel is a powerful feature that allows users to access historical data within a defined period. It

enables two key capabilities:

B. Restoring data-related objects that have been deleted within the past 90 days: Time Travel can be used to restore tables, schemas, and databases that have been accidentally or intentionally deleted within the Time Travel retention period.

C. Conducting point-in-time analysis for BI reporting: It allows users to query historical data as it appeared at a specific point in time within the Time Travel retention period, which is crucial for business intelligence and reporting purposes.

While Time Travel does allow querying of past data, it is limited to the retention period set for the Snowflake account, which is typically 1 day for standard accounts and can be extended up to 90 days for enterprise accounts. It does not enable querying or restoring objects created or deleted beyond the retention period, nor does it provide analysis over all periods of time.

References:

[Snowflake Documentation on Time Travel](#)

[SnowPro® Core Certification Study Guide](#)

Question: 31

Which statement about billing applies to Snowflake credits?

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

Answer: D

Explanation:

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

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

Question: 32

What Snowflake features allow virtual warehouses to handle high concurrency workloads? (Select TWO)

- A. The ability to scale up warehouses
- B. The use of warehouse auto scaling
- C. The ability to resize warehouses
- D. Use of multi-clustered warehouses
- E. The use of warehouse indexing

Answer: B, D

Explanation:

Snowflake's architecture is designed to handle high concurrency workloads through several features, two of which are particularly effective:

8. The use of warehouse auto scaling: This feature allows Snowflake to automatically adjust the compute resources allocated to a virtual warehouse in response to the workload. If there is an increase in concurrent queries, Snowflake can scale up the resources to maintain performance.

D. Use of multi-clustered warehouses: Multi-clustered warehouses enable Snowflake to run multiple clusters of compute resources simultaneously. This allows for the distribution of queries across clusters, thereby reducing the load on any single cluster and improving the system's ability to handle a high number of concurrent queries.

These features ensure that Snowflake can manage varying levels of demand without manual intervention, providing a seamless experience even during peak usage.

References:

[Snowflake Documentation on Virtual Warehouses](#)

[SnowPro® Core Certification Study Guide](#)

Question: 33

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: A, B

Explanation:

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](#)

Question: 34

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

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

Answer: D

Explanation:

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

References:

[Snowflake Documentation on Secure Data Sharing](#)

Question: 35

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

Explanation:

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](#)

Question: 36

A user unloaded a Snowflake table called mytable to an internal stage called mystage.

Which command can be used to view the list of files that has been uploaded to the staged?

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

Answer: D

Explanation:

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

References:

[Snowflake Documentation on Stages](#)

[SnowPro® Core Certification Study Guide](#)

Question: 37

Which of the following Snowflake capabilities are available in all Snowflake editions? (Select TWO)

- A. Customer-managed encryption keys through Tri-Secret Secure
- B. Automatic encryption of all data
- C. Up to 90 days of data recovery through Time Travel
- D. Object-level access control
- E. Column-level security to apply data masking policies to tables and views

Answer: B, D

Explanation:

In all Snowflake editions, two key capabilities are universally available:

B. Automatic encryption of all data: Snowflake automatically encrypts all data stored in its platform, ensuring security and compliance with various regulations. This encryption is transparent to users and does not require any configuration or management.

D. Object-level access control: Snowflake provides granular access control mechanisms that allow administrators to define permissions at the object level, including databases, schemas, tables, and views. This ensures that only authorized users can access specific data objects.

These features are part of Snowflake's commitment to security and governance, and they are included in every edition of the Snowflake Data Cloud.

References:

[Snowflake Documentation on Security Features](#)

Question: 38

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

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

Answer: A

Explanation:

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

References:

[Snowflake Documentation on Unloading Data](#)

[Snowflake SnowPro Core: Copy Into Command to Unload Rows to Files in Named Stage](#)

Question: 39

How often are encryption keys automatically rotated by Snowflake?

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

Answer: A

Explanation:

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

References:

[Understanding Encryption Key Management in Snowflake](#)

Question: 40

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

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

Answer: B, D

Explanation:

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

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

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

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

References:

[Snowflake Documentation on Semi-Structured Data Types](#)

[SnowPro® Core Certification Study Guide](#)

Question: 41

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

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

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

Answer: D

Explanation:

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

References:

[Snowflake Documentation on Snowpipe](#)

[SnowPro® Core Certification Study Guide](#)

Question: 42

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

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

Answer: C, D

Explanation:

When loading data into Snowflake, it is recommended to:

- C . Load files that are approximately 100-250 MB (or larger): This size is optimal for parallel processing and can help to maximize throughput. Smaller files can lead to overhead that outweighs the actual data processing time.
- D . Avoid using embedded characters such as commas for numeric data types: Embedded characters can cause issues during data loading as they may be interpreted incorrectly. It's best to clean the data of such characters to ensure accurate and efficient data loading.

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

References:

[Snowflake Documentation on Data Loading Considerations](#)

[COF-C02] SnowPro Core Certification Exam Study Guide

Question: 43

A user has 10 files in a stage containing new customer data. The ingest operation completes with no errors, using the following command:

```
COPY INTO my table FROM @my stage;
```

The next day the user adds 10 files to the stage so that now the stage contains a mixture of new customer data and updates to the previous data. The user did not remove the 10 original files.

If the user runs the same copy into command what will happen?

- A. All data from all of the files on the stage will be appended to the table
- B. Only data about new customers from the new files will be appended to the table
- C. The operation will fail with the error uncertain files in stage.
- D. All data from only the newly-added files will be appended to the table.

Answer: A

Explanation:

When the COPY INTO command is executed in Snowflake, it processes all files present in the specified stage that have not been ingested before or marked as already loaded. Since the user did not remove the original 10 files after the first load, running the same COPY INTO command again will result in all 20 files being processed. This means that the data from the original 10 files will be appended to the table again, along with the data from the new 10 files, potentially leading to duplicate records for the original data set.

References:

[Snowflake Documentation on Data Loading](#)

[SnowPro® Core Certification Study Guide](#)

Question: 44

A user has unloaded data from Snowflake to a stage

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

- A. list @file stage
- B. show @file stage

- C. view @file stage
- D. verify @file stage

Answer: A

Explanation:

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

References:

[Snowflake Documentation on Stages](#)

[SnowPro® Core Certification Study Guide](#)

Question: 45

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

Explanation:

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 scenarios1.](#)

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 properties2.](#)

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:

Question: 46

Which data types does Snowflake support when querying semi-structured data? (Select TWO)

- A. VARIANT
- B. ARRAY
- C. VARCHAR
- D. XML
- E. BLOB

Answer: A, B

Explanation:

Snowflake supports querying semi-structured data using specific data types that are capable of handling the flexibility and structure of such data. The data types supported for this purpose are:

A. VARIANT: This is a universal data type that can store values of any other type, including structured and semi-structured types. [It is particularly useful for handling JSON, Avro, ORC, Parquet, and XML data formats](#)¹.

B. ARRAY: [An array is a list of elements that can be of any data type, including VARIANT, and is used to handle semi-structured data that is naturally represented as a list](#)¹.

These data types are part of Snowflake's built-in support for semi-structured data, allowing for the storage, querying, and analysis of data that does not fit into the traditional row-column format.

References:

[Snowflake Documentation on Semi-Structured Data](#)

[COF-C02] SnowPro Core Certification Exam Study Guide

Question: 47

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

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

- C. All Snowflake accounts must be hosted in the same cloud vendor and region
- D. Each Snowflake account can be hosted in a different cloud vendor, but must be in the same region.

Answer: A

Explanation:

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

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

Question: 48

A user is loading JSON documents composed of a huge array containing multiple records into Snowflake. The user enables the strip_outer_array file format option

What does the STRIP_OUTER_ARRAY file format do?

- A. It removes the last element of the outer array.
- B. It removes the outer array structure and loads the records into separate table rows,
- C. It removes the trailing spaces in the last element of the outer array and loads the records into separate table columns
- D. It removes the NULL elements from the JSON object eliminating invalid data and enables the ability to load the records

Answer: B

Explanation:

The STRIP_OUTER_ARRAY file format option in Snowflake is used when loading JSON documents that are composed of a large array containing multiple records. When this option is enabled, it removes the outer array structure, which allows each record within the array to be loaded as a separate row in the table. [This is particularly useful for efficiently loading JSON data that is structured as an array of records1.](#)

References:

[Snowflake Documentation on JSON File Format](#)

[COF-C02] SnowPro Core Certification Exam Study Guide

Question: 49

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

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

Answer: C

Explanation:

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

References:

[Snowflake Documentation on Storage Costs for Time Travel and Fail-safe](#)

Question: 50

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

Explanation:

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-C02] SnowPro Core Certification Exam Study Guide

[Section 1.3 - SnowPro Core Certification Study Guide¹](#)

Question: 51

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

Explanation:

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

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Secure Data Sharing1](#)

Question: 52

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

Explanation:

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 Performance1](#)

Question: 53

Which Snowflake partner specializes in data catalog solutions?

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

Answer: A

Explanation:

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

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake's official documentation and partner listings](#)

Question: 54

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

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

Answer: B

Explanation:

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](#)

Question: 55

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

Explanation:

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](#)

Question: 56

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

- A. account usage . query history view
- B. account usage.query history archive View
- C. information schema . cruery_history view
- D. information schema - query history_by_session view

Answer: B

Explanation:

To find information about query activity from 90 days ago, a Snowflake user should use

the account_usage.query_history_archive view. This view is designed to provide access to historical query data beyond the default 14-day retention period found in the standard query_history view. It allows users to analyze and audit past query activities for up to 365 days after the date of execution, which includes the 90-day period mentioned.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Account Usage Schema1](#)

Question: 57

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

Explanation:

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](#)

Question: 58

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

Explanation:

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-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Network Policies1](#)

[Section 1.3 - SnowPro Core Certification Study Guide](#)

Question: 59

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

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

Answer: D

Explanation:

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

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Data Loading](#)

Question: 60

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

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

- D. When dropping an internal stage, the files are deleted with the stage and the files are not recoverable
- E. When dropping an internal stage, only selected files are deleted with the stage and are not recoverable

Answer: A, D

Explanation:

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

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

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Stages](#)

Question: 61

How long is Snowpipe data load history retained?

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

Answer: C

Explanation:

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

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Snowpipe1](#)

Question: 62

In which use cases does Snowflake apply egress charges?

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

Answer: C

Explanation:

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

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Data Replication and Egress Charges1](#)

Question: 63

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

- A. ROLES
- B. POLICY_REFERENCES
- C. QUERY_HISTORY
- D. RESOURCE_MONITORS
- E. ACCESS_HISTORY

Answer: B, E

Explanation:

To evaluate the details of dynamic data masking,

the POLICY_REFERENCES and ACCESS_HISTORY views in the account_usage schema are used.

The POLICY_REFERENCES view provides information about the objects to which a masking policy is applied, and the ACCESS_HISTORY view contains details about access to the masked data, which can be used to audit and verify the application of dynamic data masking policies.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Dynamic Data Masking1](#)

Question: 64

Query compilation occurs in which architecture layer of the Snowflake Cloud Data Platform?

- A. Compute layer
- B. Storage layer
- C. Cloud infrastructure layer
- D. Cloud services layer

Answer: D

Explanation:

Query compilation in Snowflake occurs in the Cloud Services layer. This layer is responsible for coordinating and managing all aspects of the Snowflake service, including authentication, infrastructure management, metadata management, query parsing and optimization, and security. By handling these tasks, the Cloud Services layer enables the Compute layer to focus on executing queries, while the Storage layer is dedicated to persistently storing data.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Snowflake Architecture1](#)

Question: 65

Which is the MINIMUM required Snowflake edition that a user must have if they want to use AWS/Azure Privatelink or Google Cloud Private Service Connect?

- A. Standard
- B. Premium
- C. Enterprise
- D. Business Critical

Answer: D

Explanation:

<https://docs.snowflake.com/en/user-guide/admin-security-privatelink.html>

Question: 66

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

Explanation:

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 Profiling1](#)

Question: 67

A marketing co-worker has requested the ability to change a warehouse size on their medium virtual warehouse called mktg WH.

Which of the following statements will accommodate this request?

- A. ALLOW RESIZE ON WAREHOUSE MKTG WH TO USER MKTG LEAD;
- B. GRANT MODIFY ON WAREHOUSE MKTG WH TO ROLE MARKETING;

- C. GRANT MODIFY ON WAREHOUSE MKTG WH TO USER MKTG LEAD;
- D. GRANT OPERATE ON WAREHOUSE MKTG WH TO ROLE MARKETING;

Answer: B

Explanation:

The correct statement to accommodate the request for a marketing co-worker to change the size of their medium virtual warehouse called mktg WH is to grant the MODIFY privilege on the warehouse to the ROLE MARKETING. This privilege allows the role to change the warehouse size among other properties.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Access Control Privileges1](#)

Question: 68

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

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

Answer: D

Explanation:

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

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Query Profile1](#)

[Snowpro Core Certification Exam Flashcards2](#)

Question: 69

Which stage type can be altered and dropped?

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

Answer: B

Explanation:

External stages can be altered and dropped in Snowflake. An external stage points to an external location, such as an S3 bucket, where data files are stored. Users can modify the stage's definition or drop it entirely if it's no longer needed. This is in contrast to table stages, which are tied to specific tables and cannot be altered or dropped independently.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Stages1](#)

Question: 70

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

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

Answer: A

Explanation:

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

References:

SnowPro Core Certification Exam Study Guide

Snowflake Documentation on SnowSQL

<https://docs.snowflake.com/en/user-guide/snowsql-use.html>

Question: 71

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

Explanation:

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 Snowpipe1](#)

Question: 72

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

Explanation:

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

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation12](#)

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

Question: 73

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

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

Answer: B, D

Explanation:

The Snowflake storage layer is responsible for storing data in an optimized, compressed, columnar format. This includes micro-partitions, which are the fundamental storage units that contain the actual data stored in Snowflake. Additionally, persisted query results, which are the results of queries that have been materialized and stored for future use, are also kept within this layer. [This design allows for efficient data retrieval and management within the Snowflake architecture1.](#)

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Key Concepts & Architecture | Snowflake Documentation2](#)

Question: 74

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

Explanation:

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 usage1](#)

Question: 75

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

Explanation:

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](#)

Question: 76

What is a responsibility of Snowflake's virtual warehouses?

- A. Infrastructure management
- B. Metadata management
- C. Query execution
- D. Query parsing and optimization
- E. Management of the storage layer

Answer: C

Explanation:

The primary responsibility of Snowflake's virtual warehouses is to execute queries. Virtual warehouses are one of the key components of Snowflake's architecture, providing the compute power required to perform data processing tasks such as running SQL queries, performing joins, aggregations, and other data manipulations.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Virtual Warehouses1](#)

Question: 77

Which of the following compute resources or features are managed by Snowflake? (Select TWO).

- A. Execute a COPY command
- B. Updating data
- C. Snowpipe
- D. AUTOMATIC CLUSTERING
- E. Scaling up a warehouse

Answer: C, E

Explanation:

Snowflake manages various compute resources and features, including Snowpipe and the ability to scale up a warehouse. Snowpipe is Snowflake's continuous data ingestion service that allows users to load data as soon as it becomes available. Scaling up a warehouse refers to increasing the compute resources allocated to a virtual warehouse to handle larger workloads or improve performance.

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Snowpipe and Virtual Warehouses1](#)

Question: 78

What happens when a virtual warehouse is resized?

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

Answer: A

Explanation:

When a virtual warehouse in Snowflake is resized, specifically when it is increased in size, the additional compute resources become immediately available to all running and queued queries. This means that the performance of these queries can improve due to the increased resources. [Conversely, when the size of a warehouse is reduced, the compute resources are not removed until they are no longer being used by any current operations1.](#)

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Virtual Warehouses2](#)

Question: 79

A developer is granted ownership of a table that has a masking policy. The developer's role is not able to see the masked data. Will the developer be able to modify the table to read the masked data?

- A. Yes, because a table owner has full control and can unset masking policies.
- B. Yes, because masking policies only apply to cloned tables.
- C. No, because masking policies must always reference specific access roles.
- D. No, because ownership of a table does not include the ability to change masking policies

Answer: D

Explanation:

Even if a developer is granted ownership of a table with a masking policy, they will not be able to modify the table to read the masked data if their role does not have the necessary permissions. Ownership of a table does not automatically confer the ability to alter masking policies, which are designed to protect sensitive data. [Masking policies are applied at the schema level and require specific privileges to modify¹²](#).

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Masking Policies](#)

Question: 80

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

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

Answer: B

Explanation:

Clustering keys in Snowflake work by sorting the designated columns over time. This process is done in the background and does not block data manipulation language (DML) operations, allowing for normal database operations to continue without interruption. [The purpose of clustering keys is to organize the data within micro-partitions to optimize query performance¹](#).

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Clustering¹](#)

Question: 81

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

Explanation:

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 capabilities2.](#)

References:

[COF-C02] SnowPro Core Certification Exam Study Guide

[Snowflake Documentation on Machine Learning & Data Science Partners](#)

<https://docs.snowflake.com/en/user-guide/ecosystem-analytics.html>

Question: 82

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

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

Answer: D

Explanation:

In Snowflake, when the account is located on Microsoft Azure, a valid source for an external stage can be an Azure container or a folder path within an Azure container. This includes Azure Blob storage which is accessible via the azure:// endpoint. A Windows server file share on Azure, if configured properly, can be a valid source for staging data files for Snowflake. [Options A, B, and C are not supported as direct sources for an external stage in Snowflake on Azure12.](#) References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 83

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

- A. Variant

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

Answer: D

Explanation:

Snowflake supports two geospatial data types: GEOGRAPHY and GEOMETRY. The GEOGRAPHY data type is used to store geospatial data that models the Earth as a perfect sphere, which is suitable for global geospatial data. This data type follows the WGS 84 standard and is used for storing points, lines, and polygons on the Earth's surface. The GEOMETRY data type, on the other hand, represents features in a planar (Euclidean, Cartesian) coordinate system and is typically used for local spatial reference systems. [Since the question specifically asks about geospatial data, which commonly refers to Earth-related spatial data, the correct answer is GEOGRAPHY3.](#) References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 84

What can be used to view warehouse usage over time? (Select Two).

- A. The load HISTORY view
- B. The Query history view
- C. The show warehouses command
- D. The WAREHOUSE_METERING HISTORY View
- E. The billing and usage tab in the Snowflake web UI

Answer: B, D

Explanation:

To view warehouse usage over time, the Query history view and the WAREHOUSE_METERING HISTORY View can be utilized. [The Query history view allows users to monitor the performance of their queries and the load on their warehouses over a specified period1.](#) [The WAREHOUSE_METERING HISTORY View provides detailed information about the workload on a warehouse within a specified date range, including average running and queued loads2.](#) References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 85

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

- A. Cluster keys

B. Time Travel

C. Fail-safe

D. Cloning

Answer: B

Explanation:

[Snowflake's Time Travel feature is used for both querying historical data in tables and restoring and cloning historical data in databases, schemas, and tables3](#). It allows users to access historical data within a defined period (1 day by default, up to 90 days for Snowflake Enterprise Edition) and is a key feature for data recovery and management.

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

Question: 86

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

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

A. Okta Verify

B. Duo Mobile

C. Microsoft Authenticator

D. Google Authenticator

Answer: B

Explanation:

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

Snowflake. [This service is managed entirely by Snowflake, and users do not need to sign up separately with Duo1](#).

Question: 87

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

A. Pipes

B. Streams

C. Tasks

D. Procedures

Answer: B

Explanation:

In Snowflake, Streams are the objects that track Data Manipulation Language (DML) changes made to tables, such as inserts, updates, and deletes. Streams record these changes along with metadata about each change, enabling actions to be taken using the changed data. [This process is known as change data capture \(CDC\)2.](#)

Question: 88

What tasks can be completed using the copy command? (Select TWO)

- A. Columns can be aggregated
- B. Columns can be joined with an existing table
- C. Columns can be reordered
- D. Columns can be omitted
- E. Data can be loaded without the need to spin up a virtual warehouse

Answer: C, D

Explanation:

The COPY command in Snowflake allows for the reordering of columns as they are loaded into a table, and it also permits the omission of columns from the source file during the load process. This provides flexibility in handling the schema of the data being ingested. References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 89

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

Explanation:

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>

Question: 90

What SQL command would be used to view all roles that were granted to user.1?

- A. show grants to user USER1;
- B. show grants of user USER1;
- C. describe user USER1;
- D. show grants on user USER1;

Answer: A

Explanation:

The correct command to view all roles granted to a specific user in Snowflake is SHOW GRANTS TO USER <user_name>. This command lists all access control privileges that have been explicitly granted to the specified user.

References: [SHOW GRANTS | Snowflake Documentation](#)

Question: 91

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

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

Answer: C

Explanation:

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

Question: 92

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

- A. ACCOUNTADMIN

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

Answer: A

Explanation:

The ACCOUNTADMIN role is required for a user to create and manage accounts in Snowflake. [This role has the highest level of privileges and is responsible for managing all aspects of the Snowflake account, including the ability to create and manage other user accounts1.](#)

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

Question: 93

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

Explanation:

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 analysis2](#)

Question: 94

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

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

Answer: C, E

Explanation:

The result set cache in Snowflake is invalidated and no longer available when the underlying data of the query results has changed, ensuring that queries return the most current data. [Additionally, the cache expires after 24 hours to maintain the efficiency and accuracy of data retrieval1.](#)

Question: 95

Which of the following describes external functions in Snowflake?

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

Answer: A

Explanation:

External functions in Snowflake are a special type of User-Defined Function (UDF) that call code executed outside of Snowflake, typically through a remote service. [Unlike traditional UDFs, external functions do not contain SQL code within Snowflake; instead, they interact with external services to process data2.](#)

<https://docs.snowflake.com/en/sql-reference/external-functions.html#:~:text=External%20functions%20are%20user%2Ddefined,code%20running%20outside%20of%20Snowflake.>

Question: 96

What are ways to create and manage data shares in Snowflake? (Select TWO)

- A. Through the Snowflake web interface (UI)
- B. Through the DATA_SHARE=TRUE parameter
- C. Through SQL commands
- D. Through the enable share=true parameter
- E. Using the CREATE SHARE AS SELECT * TABLE command

Answer: A, C

Explanation:

Data shares in Snowflake can be created and managed through the Snowflake web interface, which provides a user-friendly graphical interface for various operations. [Additionally, SQL commands can be used to perform these tasks programmatically, offering flexibility and automation capabilities123.](#)

Question: 97

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

A. `SELECT LAST_SUCCESS_LOGIN, LOGIN_NAME`

`FROM ACCOUNT_USAGE.USERS;`

B. `SELECT EVENT_TIMESTAMP, USER_NAME`

`FROM table(information_schema.login_history_by_user())`

C. `SELECT EVENT_TIMESTAMP, USER_NAME`

`FROM ACCOUNT_USAGE.ACCESS_HISTORY;`

D. `SELECT EVENT_TIMESTAMP, USER_NAME`

`FROM ACCOUNT_USAGE.LOGIN_HISTORY;`

Answer: D

Explanation:

To generate a report listing all Snowflake logins within the last 90 days, the `ACCOUNT_USAGE.LOGIN_HISTORY` view should be used. [This view provides information about login attempts, including successful and unsuccessful logins, and is suitable for security audits4.](#)

Question: 98

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

A. ORC

B. XML

C. Avro

D. Parquet

E. JSON

Answer: D, E

Explanation:

Semi-structured JSON, Parquet

Snowflake supports unloading data in several semi-structured file formats, including Parquet and JSON. [These formats allow for efficient storage and querying of semi-structured data, which can be loaded directly into Snowflake tables without requiring a predefined schema12.](#)

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

Question: 99

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

Explanation:

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 services3.](#)

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

Topic 2, Exam pool B

Question: 100

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

How can this worksheet be shared?

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

Answer: C

Explanation:

Worksheets in Snowsight can be shared directly with other Snowflake users within the same account. [This feature allows for collaboration and sharing of SQL queries or Python code, as well as other data manipulation tasks1.](#)

Question: 101

What is the purpose of multi-cluster virtual warehouses?

- A. To create separate data warehouses to increase query optimization
- B. To allow users the ability to choose the type of compute nodes that make up a virtual warehouse cluster
- C. To eliminate or reduce Queuing of concurrent queries
- D. To allow the warehouse to resize automatically

Answer: C

Explanation:

Multi-cluster virtual warehouses in Snowflake are designed to manage user and query concurrency needs. [They allow for the allocation of additional clusters of compute resources, either statically or dynamically, to handle increased loads and reduce or eliminate the queuing of concurrent queries2.](#)

<https://docs.snowflake.com/en/user-guide/warehouses-multicuster.html#:~:text=Multi%2Dcluster%20warehouses%20enable%20you,during%20peak%20and%20off%20hours.>

Question: 102

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

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

F. Snowflake data and services are deployed in at least three availability zones within a cloud provider's region

Answer: D, E, F

Explanation:

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

Question: 103

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

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

Answer: A, D

Explanation:

In Snowflake, both materialized views and transient tables will incur storage charges because they store data. They will also incur compute charges when queries are run against them, as compute resources are used to process the queries. References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 104

A user is preparing to load data from an external stage

Which practice will provide the MOST efficient loading performance?

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

Answer: A

Explanation:

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

Question: 105

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

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

Answer: B

Explanation:

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

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

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

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

Question: 106

Which command sets the Virtual Warehouse for a session?

- A. COPY WAREHOUSE FROM <<config file>>;
- B. SET WAREHOUSE = <<warehouse name>>;
- C. USE WAREHOUSE <<warehouse name>>;
- D. USE VIRTUAL_WAREHOUSE <<warehouse name>>;

Answer: C

Explanation:

The command USE WAREHOUSE <<warehouse name>>; is used to set the virtual warehouse for the current session in Snowflake. [This command specifies which virtual warehouse to use for executing queries in that session1.](#)

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

Question: 107

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

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

Answer: A

Explanation:

When a pipe is recreated using the CREATE OR REPLACE PIPE command, the load history of the pipe is reset. [This means that Snowpipe will consider all files in the stage as new and will attempt to load them, even if they were loaded previously by the old pipe2.](#)

Question: 108

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

- A. True
- B. False

Answer: B

Explanation:

[Snowpipe via REST API can reference both named internal stages within Snowflake and external stages, such as Amazon S3, Google Cloud Storage, or Microsoft Azure1](#). This means that Snowpipe is not limited to only external stages as a source for data loading.

References = [\[COF-C02\] SnowPro Core Certification Exam Study Guide, Snowflake Documentation1](#)

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

Question: 109

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

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

Answer: A, C, D

Explanation:

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

Question: 110

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

- A. Scaling out adds clusters of the same size to a virtual warehouse to handle more concurrent queries.
- B. Scaling out adds clusters of varying sizes to a virtual warehouse.

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

Answer: A, E

Explanation:

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

Question: 111

What is the minimum Snowflake edition that has column-level security enabled?

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

Answer: B

Explanation:

[Column-level security, which allows for the application of masking policies to columns in tables or views, is available starting from the Enterprise edition of Snowflake1.](#)

[References = \[COF-C02\] SnowPro Core Certification Exam Study Guide, Snowflake Documentation1](#)

Question: 112

When cloning a database, what is cloned with the database? (Choose two.)

- A. Privileges on the database
- B. Existing child objects within the database
- C. Future child objects within the database
- D. Privileges on the schemas within the database
- E. Only schemas and tables within the database

Answer: A, B

Explanation:

When cloning a database in Snowflake, the clone includes all privileges on the database as well as existing child objects within the database, such as schemas, tables, views, etc. [However, it does not include future child objects or privileges on schemas within the database2.](#)

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

Question: 113

Which of the following describes the Snowflake Cloud Services layer?

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

Answer: A

Explanation:

[The Snowflake Cloud Services layer is a collection of services that coordinate activities across Snowflake, tying together all the different components to process user requests, from login to query dispatch1.](#)

[References = \[COF-C02\] SnowPro Core Certification Exam Study Guide, Snowflake Documentation1](#)

Question: 114

What is the maximum total Continuous Data Protection (CDP) charges incurred for a temporary table?

- A. 30 days
- B. 7 days
- C. 48 hours
- D. 24 hours

Answer: D

Explanation:

[For a temporary table, the maximum total Continuous Data Protection \(CDP\) charges incurred are for the duration of the session in which the table was created, which does not exceed 24 hours2.](#)

[References = \[COF-C02\] SnowPro Core Certification Exam Study Guide, Snowflake Documentation2](#)

Question: 115

What type of query benefits the MOST from search optimization?

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

Answer: C

Explanation:

Search optimization in Snowflake is designed to improve the performance of queries that are selective and involve point lookup operations using equality and IN predicates. [It is particularly beneficial for queries that access columns with a high number of distinct values1.](#)

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

Question: 116

Which of the following are characteristics of Snowflake virtual warehouses? (Choose two.)

- A. Auto-resume applies only to the last warehouse that was started in a multi-cluster warehouse.
- B. The ability to auto-suspend a warehouse is only available in the Enterprise edition or above.
- C. SnowSQL supports both a configuration file and a command line option for specifying a default warehouse.
- D. A user cannot specify a default warehouse when using the ODBC driver.
- E. The default virtual warehouse size can be changed at any time.

Answer: C, E

Explanation:

Snowflake virtual warehouses support a configuration file and command line options in SnowSQL to specify a default warehouse, which is characteristic C. Additionally, the size of a virtual warehouse can be changed at any time, which is characteristic E. [These features provide flexibility and ease of use in managing compute resources2.](#)

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

Question: 117

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: D

Explanation:

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 formats¹.](#)

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

Question: 118

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

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

Answer: C

Explanation:

Snowflake's architecture is described as a multi-cluster, shared data architecture. [This design combines the simplicity of a shared-disk architecture with the performance and scale-out benefits of a shared-nothing architecture, using a central repository accessible from all compute nodes².](#)

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

Question: 119

Which of the following is a data tokenization integration partner?

- A. Protegrity
- B. Tableau

C. DBeaver

D. SAP

Answer: A

Explanation:

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

Question: 120

What versions of Snowflake should be used to manage compliance with Personal Identifiable Information (PII) requirements? (Choose two.)

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

Answer: B, C

Explanation:

To manage compliance with Personal Identifiable Information (PII) requirements, the Virtual Private Snowflake and Business Critical Editions of Snowflake should be used. [These editions provide advanced security features necessary for handling sensitive data](#)

Question: 121

What are supported file formats for unloading data from Snowflake? (Choose three.)

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

F. CSV

Answer: B, C, F

Explanation:

The supported file formats for unloading data from Snowflake include JSON, Parquet, and CSV. [These formats are commonly used for their flexibility and compatibility with various data processing tools](#)

Question: 122

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

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

Answer: B, C

Explanation:

[The Snowflake cloud services layer is responsible for tasks such as authentication and access control, ensuring secure access to the platform, and metadata management, which involves organizing and maintaining information about the data stored in Snowflake](#)

Question: 123

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

Explanation:

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 accounts1.](#)

Question: 124

When loading data into Snowflake via Snowpipe what is the compressed file size recommendation?

- A. 10-50 MB
- B. 100-250 MB
- C. 300-500 MB
- D. 1000-1500 MB

Answer: B

Explanation:

For loading data into Snowflake via Snowpipe, the recommended compressed file size is between 100-250 MB. [This size range is optimal for balancing the performance of parallel processing and minimizing the overhead associated with handling many small files2.](#)

Question: 125

Which Snowflake feature allows a user to substitute a randomly generated identifier for sensitive data, in order to prevent unauthorized users access to the data, before loading it into Snowflake?

- A. External Tokenization
- B. External Tables
- C. Materialized Views
- D. User-Defined Table Functions (UDTF)

Answer: A

Explanation:

The feature in Snowflake that allows a user to substitute a randomly generated identifier for sensitive data before loading it into Snowflake is known as External Tokenization. [This process helps to secure sensitive data by ensuring that it is not exposed in its original form, thus preventing unauthorized access3.](#)

Question: 126

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

Explanation:

[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 days1.](#)

Question: 127

A running virtual warehouse is suspended.

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

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

Answer: B

Explanation:

[When a running virtual warehouse in Snowflake is suspended and then restarted, the minimum amount of time it will incur charges for is 60 seconds2.](#)

Question: 128

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: A, B, D,

Explanation:

[The responsibilities of Snowflake's Cloud Service layer include authentication \(A\), which ensures secure access to the platform; resource management \(B\), which involves allocating and managing compute resources; and query parsing and optimization \(D\), which improves the efficiency and performance of SQL query execution³.](#)

Question: 129

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

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

Answer: A

Explanation:

Temporary and transient tables in Snowflake do not have a Fail-safe period. [Once the session ends or the tables are dropped, the data is purged and not recoverable¹.](#)

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

Question: 130

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: B

Explanation:

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

Reference: <https://docs.snowflake.com/en/sql-reference/sql/get.html>

Question: 131

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

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

Answer: C

Explanation:

[Snowflake's Fail-safe provides a 7-day period during which data in a permanent table may be recoverable, but only by Snowflake Support, not by user operations3.](#)

Question: 132

A virtual warehouse is created using the following command:

```
Create warehouse my_WH with
```

```
warehouse_size = MEDIUM
```

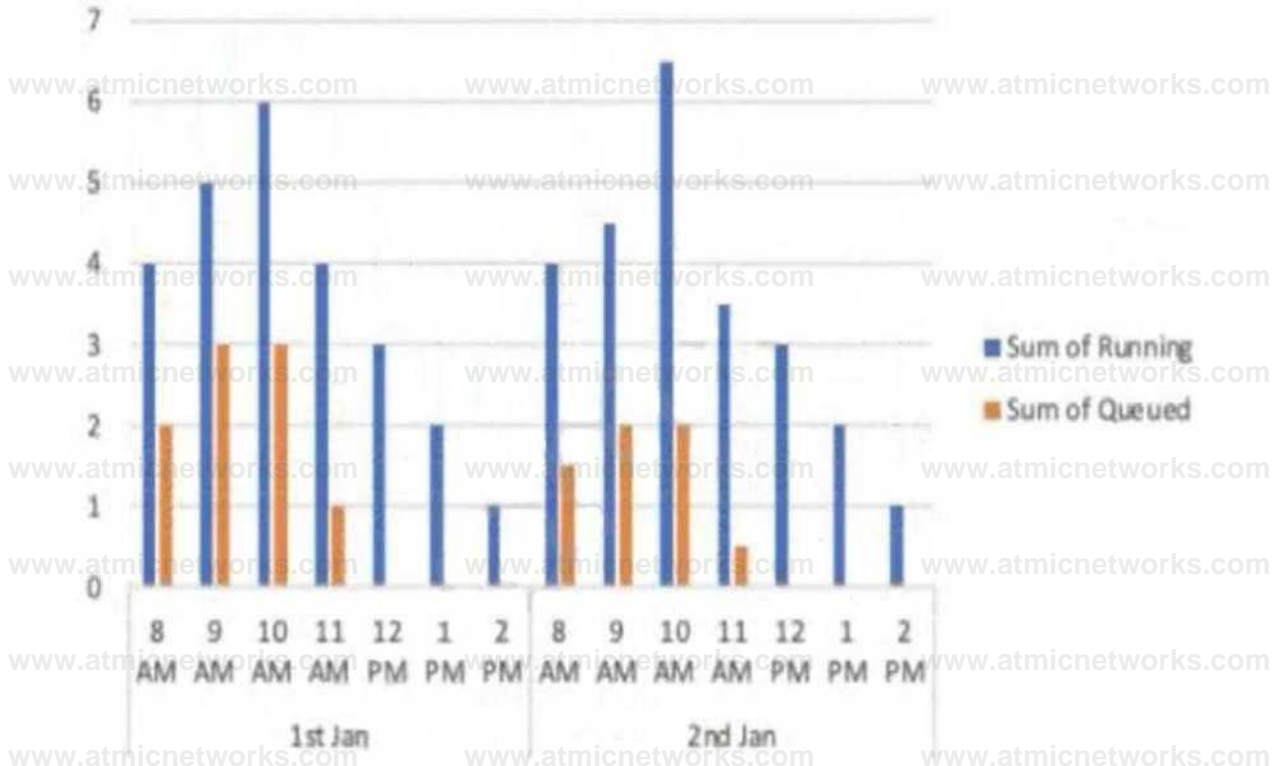
```
min_cluster_count = 1
```

```
max_cluster_count = 1
```

```
auto_suspend = 60
```

```
auto_resume = true;
```

The image below is a graphical representation of the warehouse utilization across two days.



What action should be taken to address this situation?

- Increase the warehouse size from Medium to 2XL.
- Increase the value for the parameter MAX_CONCURRENCY_LEVEL.
- Configure the warehouse to a multi-cluster warehouse.
- Lower the value of the parameter STATEMENT_QUEUED_TIMEOUT_IN_SECONDS.

Answer: C

Explanation:

The graphical representation of warehouse utilization indicates periods of significant queuing, suggesting that the current single cluster cannot efficiently handle all incoming queries. [Configuring the warehouse to a multi-cluster warehouse will distribute the load among multiple clusters, reducing queuing times and improving overall performance1.](#)

[References = Snowflake Documentation on Multi-cluster Warehouses1](#)

Question: 133

Which minimum Snowflake edition allows for a dedicated metadata store?

- Standard
- Enterprise

C. Business Critical

D. Virtual Private Snowflake

Answer: B

Explanation:

[The Enterprise edition of Snowflake allows for a dedicated metadata store, providing additional features designed for large-scale enterprises](#)

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

Question: 134

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

A. Role

B. Schema

C. User

D. Database

E. Account

F. Tables

Answer: C, E

Explanation:

[Network policies in Snowflake can be set at the user level and at the account level.](#)

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

Question: 135

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

A. Default Time Travel Retention is set to 0 days.

Maximum Time Travel Retention is 30 days.

Fail Safe retention time is 1 day.

B. Default Time Travel Retention is set to 1 day.

Maximum Time Travel Retention is 365 days.

Fail Safe retention time is 7 days.

C. Default Time Travel Retention is set to 0 days.

Maximum Time Travel Retention is 90 days.

Fail Safe retention time is 7 days.

D. Default Time Travel Retention is set to 1 day.

Maximum Time Travel Retention is 90 days.

Fail Safe retention time is 7 days.

E. Default Time Travel Retention is set to 7 days.

Maximum Time Travel Retention is 1 day.

Fail Safe retention time is 90 days.

F. Default Time Travel Retention is set to 90 days.

Maximum Time Travel Retention is 7 days.

Fail Safe retention time is 356 days.

Answer: D

Explanation:

[In the Snowflake Enterprise Edition, the default Time Travel retention is set to 1 day, the maximum Time Travel retention can be set up to 90 days, and the Fail-safe retention time is 7 days.](#)

Question: 136

Which of the following objects are contained within a schema? (Choose two.)

A. Role

B. Stream

C. Warehouse

D. External table

E. User

F. Share

Answer: B, D

Explanation:

In Snowflake, a schema is a logical grouping of database objects, which can include streams and external tables. A stream is an object that allows users to query data that has changed in specified tables or views, and an external table is a table that references data stored outside of Snowflake. Roles, warehouses, users, and shares are not contained within a schema. References: [SHOW OBJECTS, Database, Schema, & Share DDL](#)

Question: 137

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

Explanation:

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

Question: 138

A table needs to be loaded. The input data is in JSON format and is a concatenation of multiple JSON documents. The file size is 3 GB. A warehouse size small is being used. The following COPY INTO command was executed:

```
COPY INTO SAMPLE FROM @~/SAMPLE.JSON (TYPE=JSON)
```

The load failed with this error:

```
Max LOB size (16777216) exceeded, actual size of parsed column is 17894470.
```

How can this issue be resolved?

- A. Compress the file and load the compressed file.
- B. Split the file into multiple files in the recommended size range (100 MB - 250 MB).
- C. Use a larger-sized warehouse.
- D. Set STRIP_OUTER_ARRAY=TRUE in the COPY INTO command.

Answer: B

Explanation:

The error “Max LOB size (16777216) exceeded” indicates that the size of the parsed column exceeds the maximum size allowed for a single column value in Snowflake, which is 16 MB. To resolve this issue, the file should be split into multiple smaller files that are within the recommended size range of 100 MB to 250 MB. This will ensure that each JSON document within the files is smaller than the maximum LOB size allowed. Compressing the file, using a larger-sized warehouse, or setting STRIP_OUTER_ARRAY=TRUE will not resolve the issue of the column size exceeding the maximum allowed. References: [COPY INTO Error during Structured Data Load: “Max LOB size \(16777216\) exceeded...”](#)

Question: 139

Which of the following describes a Snowflake stored procedure?

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

Answer: D

Explanation:

Snowflake stored procedures can be created to execute with the privileges of the role that owns the procedure (owner's rights) or with the privileges of the role that calls the procedure (caller's rights). [This allows for flexibility in managing security and access control within Snowflake1.](#)

Question: 140

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: B, D

Explanation:

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](#)[2](#).

Question: 141

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

Explanation:

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](#)[3](#).

Question: 142

How are serverless features billed?

A. Per second multiplied by an automatic sizing for the job

B. Per minute multiplied by an automatic sizing for the job, with a minimum of one minute

C. Per second multiplied by the size, as determined by the SERVERLESS_FEATURES_SIZE account parameter

D. Serverless features are not billed, unless the total cost for the month exceeds 10% of the warehouse credits, on the account

Answer: B

Explanation:

Serverless features in Snowflake are billed based on the time they are used, measured in minutes. The cost is calculated by multiplying the duration of the job by an automatic sizing determined by Snowflake, with a minimum billing increment of one minute. This means that even if a serverless feature is used for less than a minute, it will still be billed for the full minute.

Question: 143

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

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

Answer: C

Explanation:

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.

Question: 144

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: A, C

Explanation:

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-C02] SnowPro Core Certification Exam Study Guide

Question: 145

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

Explanation:

To delete staged files from a Snowflake stage, you can specify the PURGE option in the COPY INTO <table> command, which will automatically delete the files after they have been successfully loaded. [Additionally, you can use the REMOVE command after the load completes to manually delete the files from the stage¹².](#)

References = [DROP STAGE](#), [REMOVE](#)

Question: 146


Assume there is a table consisting of five micro-partitions with values ranging from A to Z.

Which diagram indicates a well-clustered table?


A.




B.



C.



D.



- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

Explanation:

A well-clustered table in Snowflake means that the data is organized in such a way that related data points are stored close to each other within the micro-partitions. This optimizes query performance by reducing the amount of scanned data. [The diagram indicated by option C shows a well-clustered table, as it likely represents a more evenly distributed range of values across the micro-partitions1.](#)

References = [Snowflake Micro-partitions & Table Clustering](#)

Question: 147

What is an advantage of using an explain plan instead of the query profiler to evaluate the performance of a query?

- A. The explain plan output is available graphically.
- B. An explain plan can be used to conduct performance analysis without executing a query.
- C. An explain plan will handle queries with temporary tables and the query profiler will not.
- D. An explain plan's output will display automatic data skew optimization information.

Answer: B

Explanation:

An explain plan is beneficial because it allows for the evaluation of how a query will be processed without the

need to actually execute the query. [This can help in understanding the query's performance implications and potential bottlenecks without consuming resources that would be used if the query were run](#)

Question: 148

Which data types are supported by Snowflake when using semi-structured data? (Choose two.)

- A. VARIANT
- B. VARRAY
- C. STRUCT
- D. ARRAY
- E. QUEUE

Answer: A, D

Explanation:

Snowflake supports the VARIANT and ARRAY data types for semi-structured data. VARIANT can store values of any other type, including OBJECT and ARRAY, making it suitable for semi-structured data formats like JSON.

[ARRAY is used to store an ordered list of elements](#)

Question: 149

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

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

Answer: D

Explanation:

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

Question: 150

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

Explanation:

[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 periods1.](#)

Question: 151

What is the default file size when unloading data from Snowflake using the COPY command?

- A. 5 MB
- B. 8 GB
- C. 16 MB
- D. 32 MB

Answer: C

Explanation:

The default file size when unloading data from Snowflake using the COPY command is not explicitly stated in the provided resources. [However, Snowflake documentation suggests that the file size can be specified using the MAX_FILE_SIZE option in the COPY INTO <location> command2.](#)

Question: 152

What features that are part of the Continuous Data Protection (CDP) feature set in Snowflake do not require additional configuration? (Choose two.)

- A. Row level access policies
- B. Data masking policies
- C. Data encryption
- D. Time Travel
- E. External tokenization

Answer: C, D

Explanation:

Data encryption and Time Travel are part of Snowflake's Continuous Data Protection (CDP) feature set that do not require additional configuration. [Data encryption is automatically applied to all files stored on internal stages, and Time Travel allows for querying and restoring data without any extra setup](#)

Question: 153

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

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

Answer: D

Explanation:

The Cloud Services layer in Snowflake is responsible for managing the result cache. [When a query is executed, the results are stored in this cache, and subsequent identical queries can leverage these cached results without re-executing the entire query1.](#)

Question: 154

A Snowflake Administrator needs to ensure that sensitive corporate data in Snowflake tables is not visible to end users, but is partially visible to functional managers.

How can this requirement be met?

- A. Use data encryption.
- B. Use dynamic data masking.
- C. Use secure materialized views.
- D. Revoke all roles for functional managers and end users.

Answer: B

Explanation:

Dynamic data masking is a feature in Snowflake that allows administrators to define masking policies to protect sensitive data. [It enables partial visibility of the data to certain roles, such as functional managers, while](#)

[hiding it from others, like end users](#)

Question: 155

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

Explanation:

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](#)

Question: 156

What affects whether the query results cache can be used?

- A. If the query contains a deterministic function
- B. If the virtual warehouse has been suspended
- C. If the referenced data in the table has changed
- D. If multiple users are using the same virtual warehouse

Answer: C

Explanation:

The query results cache can be used as long as the data in the table has not changed since the last time the query was run. [If the underlying data has changed, Snowflake will not use the cached results and will re-execute the query1.](#)

Question: 157

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

- A. `SELECT SUM (ORDER_AMT) FROM SALES;`

- B. SELECT AVG(ORDER_QTY) FROM SALES;
- C. SELECT MIN(ORDER_AMT) FROM SALES;
- D. SELECT ORDER_AMT * ORDER_QTY FROM SALES;

Answer: B

Explanation:

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

Question: 158

How many days is load history for Snowpipe retained?

- A. 1 day
- B. 7 days
- C. 14 days
- D. 64 days

Answer: C

Explanation:

Snowpipe retains load history for 14 days. [This allows users to view and audit the data that has been loaded into Snowflake using Snowpipe within this time frame.](#)

Question: 159

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

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

Answer: A, E

Explanation:

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

Question: 160

Which command can be used to load data files into a Snowflake stage?

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

Answer: C

Explanation:

The PUT command is used to load data files into a Snowflake stage. [This command uploads data files from a local file system to a specified stage in Snowflake](#)

Question: 161

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: C, E

Explanation:

[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](#)⁴⁵.

Question: 162

What is the maximum Time Travel retention period for a temporary Snowflake table?

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

Answer: B

Explanation:

The maximum Time Travel retention period for a temporary Snowflake table is 1 day. [This is the standard retention period for temporary tables, which allows for accessing historical data within a 24-hour window](#)

Question: 163

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

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

Answer: D

Explanation:

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

Question: 164

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: A, D, E

Explanation:

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](#)

Question: 165

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

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

Answer: D

Explanation:

In Snowflake's access control model, the default owner of an object is the role that was used to create the object. [This role has the OWNERSHIP privilege on the object and can grant access to other roles1](#)

Question: 166

What is the minimum Snowflake edition required to use Dynamic Data Masking?

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

Answer: B

Explanation:

The minimum Snowflake edition required to use Dynamic Data Masking is the Enterprise edition. [This feature is not available in the Standard edition2.](#)

Question: 167

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

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

Answer: C, E

Explanation:

The Snowflake Cloud Services layer manages various services, including authentication and metadata management. [This layer ties together all the different components of Snowflake to process user requests, manage sessions, and control access3.](#)

Question: 168

A company needs to allow some users to see Personally Identifiable Information (PII) while limiting other users from seeing the full value of the PII.

Which Snowflake feature will support this?

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

Answer: B

Explanation:

Data masking policies in Snowflake allow for the obfuscation of specific data within a field, enabling some users to see the full data while limiting others. [This feature is particularly useful for handling PII, ensuring that sensitive information is only visible to authorized users1.](#)

Question: 169

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

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

- A. view @my_stage
- B. list @my_stage
- C. show @my_stage
- D. display @my_stage

Answer: B

Explanation:

The list @my_stage command in Snowflake can be used to verify if data has been uploaded to an external stage named my_stage. [This command provides a list of files that are present in the specified stage2.](#)

Question: 170

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

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

Answer: A, E

Explanation:

The Snowflake Cloud Services layer performs a variety of tasks, including the management of metadata and the parsing and optimization of queries. [This layer is responsible for coordinating activities across Snowflake, including user session management, security, and query compilation3.](#)

Question: 171

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

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

E. The Data Consumer pays only for compute resources to query the shared data.

Answer: C, E

Explanation:

In Snowflake's data sharing model, any full Snowflake account can both provide and consume shared data.

Additionally, the data consumer pays only for the compute resources used to query the shared data. [No actual data is copied or transferred between accounts, and shared data does not take up any storage in a consumer account, so the consumer does not pay for data storage1.](#)

References = [Introduction to Secure Data Sharing | Snowflake Documentation](#)

Question: 172

The following JSON is stored in a VARIANT column called src of the CAR_SALES table:

```
"customer": [
  {
    "address": "San Francisco, CA",
    "name": "Jane Doe"

    "date": "2022-01-28",
    "dealership": "Town Auto Sales",
    "salesperson": (
      n4J»i.HrrH
    )
  }
]
```

A user needs to extract the dealership information from the JSON.

How can this be accomplished?

- A. select src:dealership from car_sales;
- B. select src.dealership from car_sales;
- C. select src:Dealership from car_sales;
- D. select dealership from car_sales;

Answer: B

Explanation:

In Snowflake, to extract a specific element from a JSON stored in a VARIANT column, the correct syntax is to use the dot notation. Therefore, the query `select src.dealership from car_sales;` will return the dealership information contained within each JSON object in the src column.

References: For a detailed explanation, please refer to the Snowflake documentation on querying SEMI-structured data.

Question: 173

Which of the following significantly improves the performance of selective point lookup queries on a table?

- A. Clustering
- B. Materialized Views
- C. Zero-copy Cloning
- D. Search Optimization Service

Answer: D

Explanation:

[The Search Optimization Service significantly improves the performance of selective point lookup queries on tables by creating and maintaining a persistent data structure called a search access path, which allows some micro-partitions to be skipped when scanning the table](#)

Question: 174

Which of the following accurately describes shares?

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

Answer: A

Explanation:

Shares in Snowflake are named objects that encapsulate all the information required to share databases, schemas, tables, secure views, and secure UDFs. [These objects can be added to a share by granting](#)

[privileges on them to the share via a database role](#)

Question: 175

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: A, D

Explanation:

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

Question: 176

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

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

Answer: B

Explanation:

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

Question: 177

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

A. 1 day

B. 7 days

C. 12 hours

D. 0 days

Answer: D

Explanation:

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

Question: 178

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

Explanation:

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](#)

Question: 179

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

Explanation:

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 view credit consumption](#)

Question: 180

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

A. SINGLE=TRUE

B. MULTIPLE=TRUE

C. MULTIPLE=FALSE

D. SINGLE=FALSE

Answer: D

Explanation:

The COPY INTO SQL command with the option SINGLE=FALSE is used to unload data into multiple files. This

option allows the data to be split into multiple files during the unload

process. References: [SnowPro Core Certification COPY INTO SQL command unload multiple files](#)

Question: 181

When cloning a database containing stored procedures and regular views, that have fully qualified table

references, which of the following will occur?

A. The cloned views and the stored procedures will reference the cloned tables in the cloned database.

B. An error will occur, as views with qualified references cannot be cloned.

C. An error will occur, as stored objects cannot be cloned.

D. The stored procedures and views will refer to tables in the source database.

Answer: A

Explanation:

When cloning a database containing stored procedures and regular views with fully qualified table references,

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

Question: 182

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

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

Answer: A

Explanation:

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

Study Guide

Question: 183

Which of the following objects can be directly restored using the UNDROP command? (Choose two.)

- A. Schema
- B. View
- C. Internal stage
- D. Table
- E. User
- F. Role

Answer: B, D

Explanation:

The UNDROP command in Snowflake can be used to directly restore Views and Tables. These objects, when dropped, are moved to a 'Recycle Bin' where they can be restored within a time limit before they are permanently deleted. References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 184

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

Explanation:

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

Question: 185

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

Explanation:

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-C02] SnowPro Core Certification Exam Study Guide

Question: 186

A user created a transient table and made several changes to it over the course of several days.

Three days after the table was created, the user would like to go back to the first version of the table.

How can this be accomplished?

- A. Use Time Travel, as long as DATA_RETENTION_TIME_IN_DAYS was set to at least 3 days.

- B. The transient table version cannot be retrieved after 24 hours.
- C. Contact Snowflake Support to have the data retrieved from Fail-safe storage.
- D. Use the FAIL_SAFE parameter for Time Travel to retrieve the data from Fail-safe storage.

Answer: A

Explanation:

To go back to the first version of a transient table created three days prior, one can use Time Travel if the DATA_RETENTION_TIME_IN_DAYS was set to at least 3 days. This allows the user to access historical data within the specified retention period. References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 187

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

Explanation:

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-C02] SnowPro Core Certification Exam Study Guide

Question: 188

Where can a user find and review the failed logins of a specific user for the past 30 days?

- A. The USERS view in ACCOUNT_USAGE
- B. The LOGIN_HISTORY view in ACCOUNT_USAGE
- C. The ACCESS_HISTORY view in ACCOUNT_USAGE
- D. The SESSIONS view in ACCOUNT_USAGE

Answer: B

Explanation:

The LOGIN_HISTORY view in the ACCOUNT_USAGE schema provides information about login attempts, including both successful and failed logins. This view can be used to review the failed login attempts of a specific user for the past 30 days. References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 189

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

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

Answer: A, C

Explanation:

Snowflake uses a Role-Based Access Control (RBAC) model to manage access to data and resources.

Additionally, Snowflake ensures that all data is encrypted, both at rest and in transit, to provide a high level of security for data stored within the platform. References: [COF-C02] SnowPro Core Certification Exam

Study Guide

Question: 190

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

1. Runs a query that takes 3 minutes to complete
2. Leaves for 15 minutes
3. Returns and runs a query that takes 10 seconds to complete
4. Manually suspends the warehouse as soon as the last query was completed

When the user returns, how much billable compute time will have been consumed?

- A. 4 minutes
- B. 10 minutes

C. 14 minutes

D. 24 minutes

Answer: C

Explanation:

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

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

Question: 191

What actions will prevent leveraging of the ResultSet cache? (Choose two.)

- A. Removing a column from the query SELECT list
- B. Stopping the virtual warehouse that the query is running against
- C. Clustering of the data used by the query
- D. Executing the RESULTS_SCAN() table function
- E. Changing a column that is not in the cached query

Answer: B, D

Explanation:

The ResultSet cache is leveraged to quickly return results for repeated queries. Actions that prevent leveraging this cache include stopping the virtual warehouse that the query is running against (B) and executing the RESULTS_SCAN() table function (D). [Stopping the warehouse clears the local disk cache, including the ResultSet cache1](#). The RESULTS_SCAN() function is used to retrieve the result of a previously executed query, which bypasses the need for the ResultSet cache.

Question: 192

Which statement is true about running tasks in Snowflake?

- A. A task can be called using a CALL statement to run a set of predefined SQL commands.
- B. A task allows a user to execute a single SQL statement/command using a predefined schedule.
- C. A task allows a user to execute a set of SQL commands on a predefined schedule.

D. A task can be executed using a SELECT statement to run a predefined SQL command.

Answer: B

Explanation:

In Snowflake, a task allows a user to execute a single SQL statement/command using a predefined schedule (B).

Tasks are used to automate the execution of SQL statements at scheduled intervals.

Question: 193

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

Explanation:

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.](#)

Question: 194

What is the following SQL command used for?

```
Select * from table(validate(t1, job_id => '_last'));
```

- A. To validate external table files in table t1 across all sessions
- B. To validate task SQL statements against table t1 in the last 14 days
- C. To validate a file for errors before it gets executed using a COPY command
- D. To return errors from the last executed COPY command into table t1 in the current session

Answer: D

Explanation:

The SQL command `Select * from table(validate(t1, job_id => '_last'));` is used to return errors from the last executed COPY command into table t1 in the current session. [It checks the results of the most recent data load](#)

[operation and provides details on any errors that occurred during that process1.](#)

Question: 195

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

Explanation:

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

Question: 196

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

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

Answer: C

Explanation:

The Snowflake Search Optimization Service is designed to support improved performance for selective point lookup queries. [These are queries that retrieve specific records from a database, often based on a unique identifier or a small set of criteria3.](#)

Question: 197

What are common issues found by using the Query Profile? (Choose two.)

- A. Identifying queries that will likely run very slowly before executing them

- B. Locating queries that consume a high amount of credits
- C. Identifying logical issues with the queries
- D. Identifying inefficient micro-partition pruning
- E. Data spilling to a local or remote disk

Answer: D, E

Explanation:

The Query Profile in Snowflake is used to identify performance issues with queries. Common issues that can be found using the Query Profile include identifying inefficient micro-partition pruning (D) and data spilling to a local or remote disk (E). [Micro-partition pruning is related to the efficiency of query execution, and data spilling occurs when the memory is insufficient, causing the query to write data to disk, which can slow down the query performance1.](#)

Question: 198

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

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

Answer: C

Explanation:

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

Question: 199

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

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

Answer: D

Explanation:

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

Question: 200

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

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

Answer: C

Explanation:

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

Question: 201

Files have been uploaded to a Snowflake internal stage. The files now need to be deleted.

Which SQL command should be used to delete the files?

- A. PURGE
- B. MODIFY

C. REMOVE

D. DELETE

Answer: C

Explanation:

The SQL command used to delete files from a Snowflake internal stage is REMOVE. [This command can be used to remove files from either an internal or external stage within Snowflake](#)

Question: 202

How should a virtual warehouse be configured if a user wants to ensure that additional multi-clusters are resumed with no delay?

- A. Configure the warehouse to a size larger than generally required
- B. Set the minimum and maximum clusters to autoscale
- C. Use the standard warehouse scaling policy
- D. Use the economy warehouse scaling policy

Answer: A

Explanation:

To ensure that additional multi-clusters are resumed with no delay, a virtual warehouse should be configured to a size larger than generally required. [This configuration allows for immediate availability of additional resources when needed, without waiting for new clusters to start up](#)

Topic 3, Exam pool C

Question: 203

Where is Snowflake metadata stored?

- A. Within the data files
- B. In the virtual warehouse layer
- C. In the cloud services layer
- D. In the remote storage layer

Answer: C

Explanation:

Snowflake's architecture is divided into three layers: database storage, query processing, and cloud services.

The metadata, which includes information about the structure of the data, the SQL operations performed, and the service-level policies, is stored in the cloud services layer. This layer acts as the brain of the Snowflake environment, managing metadata, query optimization, and transaction coordination.

Question: 204

Network policies can be applied to which of the following Snowflake objects? (Choose two.)

- A. Roles
- B. Databases
- C. Warehouses
- D. Users
- E. Accounts

Answer: D, E

Explanation:

Network policies in Snowflake can be applied to users and accounts. [These policies control inbound access to the Snowflake service and internal stages, allowing or denying access based on the originating network identifiers¹².](#)

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

Question: 205

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

- A. Configure the user to be initially disabled.
- B. Force an immediate password change.
- C. Set a default role for the user.
- D. Set the number of minutes to unlock to 15 minutes.
- E. Set the user's access to expire within a specified timeframe.

Answer: B, C

Explanation:

Question: 206

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

Explanation:

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 performance12.](#)

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

Question: 207

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

Explanation:

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 warehouse3.](#)

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

Question: 208

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

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

D. 365 Days.

Answer: A

Explanation:

Snowflake automatically rotates the Account and Table master keys when they are more than 30 days old.

[Active keys are retired, and new keys are created, ensuring robust security through frequent key changes1](#)

Question: 209

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

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

Answer: C, E

Explanation:

Clones in Snowflake have the advantage of being created almost instantly and saving space by not duplicating storage. [This is due to Snowflake's zero-copy cloning feature, which allows for the creation of object clones without the additional storage costs typically associated with data duplication23](#). Clones are independent of the original table and do not stay in sync with it, nor do they inherently have better query performance.

[However, they do inherit the time travel history from the original table at the time of cloning](#)

Question: 210

Which query profile statistics help determine if efficient pruning is occurring? (Choose two.)

- A. Bytes sent over network
- B. Percentage scanned from cache
- C. Partitions total
- D. Bytes spilled to local storage
- E. Partitions scanned

Answer: C, E

Explanation:

Efficient pruning in Snowflake is indicated by the number of partitions scanned out of the total available. [If a small percentage of partitions are scanned, it suggests that the pruning process is effectively narrowing down the data, which can lead to improved query performance](#)

Question: 211

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

- A. CHECK_JSON
- B. TO_JSON
- C. FLATTEN
- D. PARSE_JSON

Answer: C

Explanation:

The FLATTEN table function in Snowflake is used to convert semi-structured data, such as JSON or XML, into a relational format. [It expands nested arrays or objects into multiple rows, making the data suitable for relational querying.](#)

Question: 212

Which URL type allows users to access unstructured data without authenticating into Snowflake or passing an authorization token?

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

Answer: A

Explanation:

Pre-signed URLs in Snowflake allow users to access unstructured data without the need for authentication into Snowflake or passing an authorization token. [These URLs are open and can be directly accessed or downloaded by any user or application, making them ideal for business intelligence applications or reporting tools that need](#)

[to display unstructured file contents](#)

Question: 213

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

Explanation:

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](#)

Question: 214

Which of the following statements describes a schema in Snowflake?

- A. A logical grouping of objects that belongs to a single database
- B. A logical grouping of objects that belongs to multiple databases
- C. A named Snowflake object that includes all the information required to share a database
- D. A uniquely identified Snowflake account within a business entity

Answer: A

Explanation:

A schema in Snowflake is a logical grouping of database objects, such as tables and views, that belongs to a single database. [Each schema is part of a namespace in Snowflake, which is inferred from the current database and schema in use for the session](#)

Question: 215

What is a responsibility of Snowflake's virtual warehouses?

- A. Infrastructure management
- B. Metadata management
- C. Query execution

- D. Query parsing and optimization
- E. Permanent storage of micro-partitions

Answer: C

Explanation:

Snowflake's virtual warehouses are responsible for query execution. [They are clusters of compute resources that execute SQL statements, perform DML operations, and load data into tables](#)

Question: 216

Which of the following are handled by the cloud services layer of the Snowflake architecture? (Choose two.)

- A. Query execution
- B. Data loading
- C. Time Travel data
- D. Security
- E. Authentication and access control

Answer: D, E

Explanation:

[The cloud services layer of Snowflake architecture handles various aspects including security functions, authentication of user sessions, and access control, ensuring that only authorized users can access the data and services²³.](#)

Question: 217

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

Explanation:

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

Question: 218

What file formats does Snowflake support for loading semi-structured data? (Choose three.)

- A. TSV
- B. JSON
- C. PDF
- D. Avro
- E. Parquet
- F. JPEG

Answer: B, D, E

Explanation:

Snowflake supports several semi-structured data formats for loading data. [The supported formats include JSON, Avro, and Parquet](#). These formats allow for efficient storage and querying of data that does not conform to a traditional relational database schema.

Question: 219

Which Snowflake feature will allow small volumes of data to continuously load into Snowflake and will incrementally make the data available for analysis?

- A. COPY INTO
- B. CREATE PIPE
- C. INSERT INTO
- D. TABLE STREAM

Answer: B

Explanation:

The Snowflake feature that allows for small volumes of data to be continuously loaded into Snowflake and incrementally made available for analysis is Snowpipe. [Snowpipe is designed for nearreal-time data loading, enabling data to be loaded as soon as it's available in the storage layer3](#)

Question: 220

Which Snowflake objects can be shared with other Snowflake accounts? (Choose three.)

- A. Schemas
- B. Roles
- C. Secure Views
- D. Stored Procedures
- E. Tables
- F. Secure User-Defined Functions (UDFs)

Answer: A, C, F

Explanation:

In Snowflake, you can share several types of objects with other Snowflake accounts. These include schemas, secure views, and secure user-defined functions (UDFs). [Sharing these objects allows for collaboration and data access across different Snowflake accounts while maintaining security and governance controls4.](#)

Question: 221

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

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

Answer: D

Explanation:

Snowflake allows data providers to share data with consumers across different cloud platforms and regions through database replication. The data provider must replicate the database to a secondary account in the target region or cloud platform within the same organization, and then create a share to the data consumer's account. [This process ensures that the data is available in the consumer's region and on their cloud platform, facilitating seamless data sharing. References: Sharing data securely across regions and cloud platforms | Snowflake Documentation](#)

Question: 222

Which statement describes how Snowflake supports reader accounts?

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

Answer: B

Explanation:

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

Question: 223

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

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

Answer: A

Explanation:

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

Question: 224

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

Explanation:

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

Question: 225

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

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

Answer: D

Explanation:

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

Question: 226

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

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

Answer: B, D, F

Explanation:

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

Question: 227

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

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

Answer: B, C, E

Explanation:

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

Question: 228

How would a user run a multi-cluster warehouse in maximized mode?

- A. Configure the maximum clusters setting to "Maximum."
- B. Turn on the additional clusters manually after starting the warehouse.
- C. Set the minimum Clusters and maximum Clusters settings to the same value.
- D. Set the minimum clusters and maximum clusters settings to different values.

Answer: C

Explanation:

To run a multi-cluster warehouse in maximized mode, a user should set the minimum and maximum number of clusters to the same value. [This ensures that all clusters are available when the warehouse is started, providing maximum resources for query execution. References: Snowflake Documentation2.](#)

Question: 229

Data storage for individual tables can be monitored using which commands and/or objects? (Choose two.)

- A. SHOW STORAGE BY TABLE;
- B. SHOW TABLES;
- C. Information Schema -> TABLE_HISTORY
- D. Information Schema -> TABLE_FUNCTION
- E. Information Schema -> TABLE_STORAGE_METRICS

Answer: A, E

Explanation:

To monitor data storage for individual tables, the commands and objects that can be used are 'SHOW STORAGE BY TABLE;' and the Information Schema view 'TABLE_STORAGE_METRICS'. [These tools provide detailed information about the storage utilization for tables. References: Snowflake Documentation](#)

Question: 230

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

- A. 1
- B. 2

C. 3

D. 4

Answer: A

Explanation:

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

Question: 231

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`

WHERE ..

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 call stored_proc1(); call stored_proc2();`

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

Explanation:

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

Question: 232

What is the minimum Snowflake edition needed for database failover and fail-back between Snowflake accounts for business continuity and disaster recovery?

A. Standard

B. Enterprise

C. Business Critical

D. Virtual Private Snowflake

Answer: C

Explanation:

[The minimum Snowflake edition required for database failover and fail-back between Snowflake accounts for business continuity and disaster recovery is the Business Critical edition. References: Snowflake Documentation](#)

Question: 233

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

Which condition will trigger a cluster to shut-down?

A. When after 2-3 consecutive checks the system determines that the load on the most-loaded cluster could be redistributed.

B. When after 5-6 consecutive checks the system determines that the load on the most-loaded cluster could be redistributed.

C. When after 5-6 consecutive checks the system determines that the load on the least-loaded cluster could be redistributed.

D. When after 2-3 consecutive checks the system determines that the load on the least-loaded cluster could be redistributed.

Answer: D

Explanation:

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

Question: 234

What is cached during a query on a virtual warehouse?

A. All columns in a micro-partition

B. Any columns accessed during the query

C. The columns in the result set of the query

D. All rows accessed during the query

Answer: C

Explanation:

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

Question: 235

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

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

Answer: B, D

Explanation:

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

Exam Study Guide

Question: 236

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

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

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

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

Answer: B

Explanation:

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

Question: 237

Which of the following can be used when unloading data from Snowflake? (Choose two.)

- A. When unloading semi-structured data, it is recommended that the STRIP_OUTER_ARRAY option be used.
- B. Use the ENCODING file format option to change the encoding from the default UTF-8.
- C. The OBJECT_CONSTRUCT function can be used to convert relational data to semi-structured data.
- D. By using the SINGLE = TRUE parameter, a single file up to 5 GB in size can be exported to the storage layer.
- E. Use the PARSE_JSON function to ensure structured data will be unloaded into the VARIANT data type.

Answer: C, D

Explanation:

The OBJECT_CONSTRUCT function is used in Snowflake to create a JSON object from relational data, which is useful when unloading semi-structured data. The SINGLE = TRUE parameter is used when unloading data to ensure that the data is exported as a single file, which can be up to 5 GB in size. References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 238

What effect does WAIT_FOR_COMPLETION = TRUE have when running an ALTER WAREHOUSE command and changing the warehouse size?

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

Answer: D

Explanation:

The WAIT_FOR_COMPLETION = TRUE parameter in an ALTER WAREHOUSE command ensures that the command does not return until the warehouse has completed resizing. This means that the command will wait until all the necessary compute resources have been provisioned and the warehouse size has been changed.

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

Question: 239

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

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

Answer: A, C

Explanation:

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

Question: 240

Which of the following are characteristics of security in Snowflake?

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

Answer: C

Explanation:

[One of the security features of Snowflake includes the periodic rekeying of encrypted data, which is available with the Snowflake Enterprise edition and higher²](#). 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

Question: 241

User INQUISITIVE_PERSON has been granted the role DATA_SCIENCE. The role DATA_SCIENCE has privileges OWNERSHIP on the schema MARKETING of the database ANALYTICS_DW.

Which command will show all privileges granted to that schema?

- A. SHOW GRANTS ON ROLE DATA_SCIENCE
- B. SHOW GRANTS ON SCHEMA ANALYTICS_DW.MARKETING
- C. SHOW GRANTS TO USER INQUISITIVE_PERSON
- D. SHOW GRANTS OF ROLE DATA_SCIENCE

Answer: B

Explanation:

[To show all privileges granted to a specific schema, the command SHOW GRANTS ON SCHEMA <schema name> should be used³](#). In this case, it would be SHOW GRANTS ON SCHEMA ANALYTICS_DW.MARKETING. References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 242

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

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

Answer: D

Explanation:

[The bulk data load history available after a COPY statement is stored in the metadata of the pipe and is retained for 64 days¹](#).

Question: 243

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

- A. CHECK_OPTION
- B. IS_SECURE
- C. IS_UPDATEABLE
- D. TABLE_NAME

Answer: B

Explanation:

[In the Information Schema and Account Usage, the column that identifies whether a view is secure or not is IS_SECURE2.](#)

Question: 244

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

How does Snowflake recommend maximizing the load performance?

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

Answer: B

Explanation:

Snowflake's documentation recommends producing the largest files possible for data loading, as larger files reduce the number of files to process and the overhead associated with handling many small files. [This approach can maximize the load performance by leveraging Snowflake's ability to ingest large files efficiently1.](#)

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

Question: 245

If a Snowflake user decides a table should be clustered, what should be used as the cluster key?

- A. The columns that are queried in the select clause.
- B. The columns with very high cardinality.
- C. The columns with many different values.
- D. The columns most actively used in the select filters.

Answer: D

Explanation:

When deciding on a clustering key for a table, Snowflake recommends using the columns that are most actively used in the select filters. [This is because clustering by these columns can improve the performance of queries that filter on these values, leading to more efficient scans and better overall query performance2](#). References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 246

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

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

Answer: A

Explanation:

Periodic rekeying of micro-partitions is a feature that requires the Enterprise Edition of Snowflake or higher. [This feature is part of Snowflake's comprehensive approach to encryption key management, ensuring data security through regular rekeying1](#). References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 247

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

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

Answer: C

Explanation:

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

Question: 248

A materialized view should be created when which of the following occurs? (Choose two.)

- A. There is minimal cost associated with running the query.
- B. The query consumes many compute resources every time it runs.
- C. The base table gets updated frequently.
- D. The query is highly optimized and does not consume many compute resources.
- E. The results of the query do not change often and are used frequently.

Answer: B, E

Explanation:

A materialized view is beneficial when the query consumes many compute resources every time it runs (B), and when the results of the query do not change often and are used frequently (E). [This is because materialized views store pre-computed data, which can speed up query performance for workloads that are run frequently or are complex](#)

Question: 249

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

Explanation:

[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 view2.](#) This privilege is necessary before granting the SELECT privilege on a secure view to a share.

Question: 250

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

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

Answer: C

Explanation:

Automatic Clustering is a feature that allows users to control the organization of data within micropartitions in Snowflake. [By defining clustering keys, Snowflake can automatically reorganize the data in micro-partitions to optimize query performance1.](#)

Question: 251

A Snowflake user executed a query and received the results. Another user executed the same query 4 hours later. The data had not changed.

What will occur?

- A. No virtual warehouse will be used, data will be read from the result cache.
- B. No virtual warehouse will be used, data will be read from the local disk cache.
- C. The default virtual warehouse will be used to read all data.
- D. The virtual warehouse that is defined at the session level will be used to read all data.

Answer: A

Explanation:

Snowflake maintains a result cache that stores the results of every query for 24 hours. [If the same query is executed again within this time frame and the data has not changed, Snowflake will retrieve the data from the result cache instead of using a virtual warehouse to recompute the results2.](#)

Question: 252

Which statements reflect key functionalities of a Snowflake Data Exchange? (Choose two.)

- A. If an account is enrolled with a Data Exchange, it will lose its access to the Snowflake Marketplace.
- B. A Data Exchange allows groups of accounts to share data privately among the accounts.
- C. A Data Exchange allows accounts to share data with third, non-Snowflake parties.
- D. Data Exchange functionality is available by default in accounts using the Enterprise edition or higher.
- E. The sharing of data in a Data Exchange is bidirectional. An account can be a provider for some datasets and a consumer for others.

Answer: B, E

Explanation:

A Snowflake Data Exchange allows groups of accounts to share data privately among the accounts (B), and it supports bidirectional sharing, meaning an account can be both a provider and a consumer of data (E). [This facilitates secure and governed data collaboration within a selected group.](#)

Question: 253

Which database objects can be shared with the Snowflake secure data sharing feature? (Choose two.)

- A. Files
- B. External tables
- C. Secure User-Defined Functions (UDFs)
- D. Sequences
- E. Streams

Answer: B, C

Explanation:

Snowflake's secure data sharing feature allows sharing of certain database objects with other Snowflake accounts. [Among the options provided, external tables and secure UDFs can be shared](#)

Question: 254

Query parsing and compilation occurs in which architecture layer of the Snowflake Cloud Data Platform?

- A. Cloud services layer
- B. Compute layer
- C. Storage layer

D. Cloud agnostic layer

Answer: A

Explanation:

Query parsing and compilation in Snowflake occur within the cloud services layer. [This layer is responsible for various management tasks, including query compilation and optimization](#)

Question: 255

Which of the following is the Snowflake Account_Usage.Metering_History view used for?

- A. Gathering the hourly credit usage for an account
- B. Compiling an account's average cloud services cost over the previous month
- C. Summarizing the throughput of Snowpipe costs for an account
- D. Calculating the funds left on an account's contract

Answer: A

Explanation:

The Snowflake Account_Usage.Metering_History view is used to gather the hourly credit usage for an account. [This view provides details on the credits consumed by various services within Snowflake for the last 365 days1.](#)

Question: 256

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

- A. SHOW MASKING POLICIES
- B. DESCRIBE MASKING POLICY
- C. GET_DDL
- D. LIST MASKING POLICIES

Answer: C

Explanation:

The SQL command GET_DDL can be used to retrieve the CREATE definition of a masking policy in Snowflake. [This command generates the DDL statement required to recreate the masking policy](#)

Question: 257

Which statement describes pruning?

- A. The filtering or disregarding of micro-partitions that are not needed to return a query.
- B. The return of micro-partitions values that overlap with each other to reduce a query's runtime.
- C. A service that is handled by the Snowflake Cloud Services layer to optimize caching.
- D. The ability to allow the result of a query to be accessed as if it were a table.

Answer: A

Explanation:

Pruning in Snowflake refers to the process of filtering or disregarding micro-partitions that are not needed to satisfy the conditions of a query. [This optimization technique helps reduce the amount of data scanned, thereby improving query performance](#)

Question: 258

How can a user change which columns are referenced in a view?

- A. Modify the columns in the underlying table
- B. Use the ALTER VIEW command to update the view
- C. Recreate the view with the required changes
- D. Materialize the view to perform the changes

Answer: C

Explanation:

In Snowflake, to change the columns referenced in a view, the view must be recreated with the required changes. The ALTER VIEW command does not allow changing the definition of a view; it can only be used to rename a view, convert it to or from a secure view, or add, overwrite, or remove a comment for a view. [Therefore, the correct approach is to drop the existing view and create a new one with the desired column references.](#)

Question: 259

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

Explanation:

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.](#)

Question: 260

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

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

Answer: C

***E*xplanation:**

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

Question: 261

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

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

Answer: C

Explanation:

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

Question: 262

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: D

Explanation:

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

Question: 263

Which Snowflake tool would be BEST to troubleshoot network connectivity?

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

Answer: D

Explanation:

SnowCD (Snowflake Connectivity Diagnostic Tool) is the best tool provided by Snowflake for troubleshooting network connectivity issues. It helps diagnose and resolve issues related to connecting to Snowflake

services

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

[guide/snowcd.html#:~:text=SnowCD%20\(i.e.%20Snowflake%20Connectivity%20Diagnostic,their%20](https://docs.snowflake.com/en/user-guide/snowcd.html#:~:text=SnowCD%20(i.e.%20Snowflake%20Connectivity%20Diagnostic,their%20)

[network%20connection%20to%20Snowflake.](#)

Question: 264

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

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

Answer: D, E

Explanation:

The Cloud Services layer in Snowflake includes activities such as user authentication and infrastructure management. [This layer coordinates activities across Snowflake, including security enforcement, query compilation and optimization, and more](#)

Question: 265

What is the purpose of using the OBJECT_CONSTRUCT function with the COPY INTO command?

- A. Reorder the rows in a relational table and then unload the rows into a file
- B. Convert the rows in a relational table to a single VARIANT column and then unload the rows into a file.
- C. Reorder the data columns according to a target table definition and then unload the rows into the table.
- D. Convert the rows in a source file to a single variant column and then load the rows from the file to a variant table.

Answer: B

Explanation:

The OBJECT_CONSTRUCT function is used with the COPY INTO command to convert the rows in a relational table to a single VARIANT column, which can then be unloaded into a file. [This is useful for transforming table data into a semi-structured JSON format](#)

Question: 266

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

more columns?

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

Answer: A

Explanation:

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

Question: 267

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

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

Answer: D

Explanation:

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

Question: 268

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

Explanation:

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.

Question: 269

How can a data provider ensure that a data consumer is going to have access to the required objects?

- A. Enable the data sharing feature in the account and validate the view.
- B. Use the CURRENT_ROLE and CURRENT_USER functions to validate secure views.
- C. Use the CURRENT_ function to authorize users from a specific account to access rows in a base table.
- D. Set the SIMULATED DATA SHARING CONSUMER session parameter to the name of the consumer account for which access is being simulated.

Answer: A

Explanation:

To ensure a data consumer has access to the required objects, a data provider can enable the data sharing feature and validate that the consumer can access the views or tables shared with them. References: Based on general data sharing practices in cloud services as of 2021.

Question: 270

How can a user improve the performance of a single large complex query in Snowflake?

- A. Scale up the virtual warehouse.
- B. Scale out the virtual warehouse.
- C. Enable standard warehouse scaling.
- D. Enable economy warehouse scaling.

Answer: A

Explanation:

Scaling up the virtual warehouse in Snowflake involves increasing the compute resources available for a single warehouse, which can improve the performance of large and complex queries by providing more CPU and memory resources. References: Based on general cloud data warehousing knowledge as of 2021.

Question: 271

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: C, D

Explanation:

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

Question: 272

What computer language can be selected when creating User-Defined Functions (UDFs) using the Snowpark API?

- A. Swift
- B. JavaScript
- C. Python
- D. SQL

Answer: C

Explanation:

The Snowpark API allows developers to create User-Defined Functions (UDFs) in various languages, including Python, which is known for its ease of use and wide adoption in data-related tasks. References: Based on general programming and cloud data service knowledge as of 2021.

Question: 273

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

- A. 8MB

B. 16MB

C. 32MB

D. 128MB

Answer: B

Explanation:

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.

Question: 274

What is the name of the SnowSQLfile that can store connection information?

A. history

B. config

C. snowsqlcnf

D. snowsql.pubkey

Answer: B

Explanation:

The SnowSQL file that can store connection information is named 'config'. It is used to store user

credentials and connection details for easy access to Snowflake instances. References: Based on general database knowledge as of 2021.

Question: 275

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

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

A. Rename listings.

B. Delete provider profiles.

C. Modify listings properties.

D. Modify incoming listing access requests.

E. Submit listings for approval/publishing.

Answer: C, E

Explanation:

With the create data EXCHANGE listing privilege, a Snowflake user can modify the properties of listings and submit them for approval or publishing on the Data Exchange. This allows them to manage and share data sets with consumers effectively. References: Based on general data exchange practices in cloud services as of 2021.

Question: 276

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

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

Answer: A

Explanation:

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

Question: 277

Which user object property requires contacting Snowflake Support in order to set a value for it?

- A. DISABLED
- B. MINS TO BYPASS MFA
- C. MINS TO BYPASS NETWORK POLICY
- D. MINS TO UNLOCK

Answer: B

Explanation:

[The user property 'MINS TO BYPASS MFA' in Snowflake allows temporary bypass of MFA for a user, which can be set by an account administrator without contacting Snowflake Support2.](#)

Question: 278

Which file format will keep floating-point numbers from being truncated when data is unloaded?

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

Answer: D

Explanation:

[The Parquet file format is known for preserving the precision of floating-point numbers when data is unloaded, preventing truncation of the values³.](#)

Question: 279

In which Snowflake layer does Snowflake reorganize data into its internal optimized, compressed, columnar format?

- A. Cloud Services
- B. Database Storage
- C. Query Processing

D. Metadata Management

Answer: B

Explanation:

Snowflake reorganizes data into its internal optimized, compressed, columnar format in the Database Storage layer. [This process is part of how Snowflake manages data storage, ensuring efficient data retrieval and query performance](#)

Question: 280

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

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

Answer: A

Explanation:

[In the Admin area of Snowsight, users can analyze query performance, manage Snowflake warehouses, set up and view details about resource monitors, manage users and roles, and administer Snowflake accounts in their organization2.](#)

Question: 281

How can a Snowflake user optimize query performance in Snowflake? (Select TWO).

- A. Create a view.
- B. Cluster a table.
- C. Enable the search optimization service.
- D. Enable Time Travel.
- E. Index a table.

Answer: B, C

Explanation:

To optimize query performance in Snowflake, users can cluster a table, which organizes the data in a way that

minimizes the amount of data scanned during queries. [Additionally, enabling the search optimization service can improve the performance of selective point lookup queries on large tables³⁴.](#)

Question: 282

What MINIMUM privilege is required on the external stage for any role in the GET REST API to access unstructured data files using a file URL?

- A. READ
- B. OWNERSHIP
- C. USAGK
- D. WRTTF

Answer: A

Explanation:

The minimum privilege required on an external stage for any role to access unstructured data files using a file URL in the GET REST API is READ. This allows the role to retrieve or download data files from the stage.

Question: 283

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

Explanation:

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

Question: 284

Which Snowflake URL type allows users or applications to download or access files directly from Snowflake stage without authentication?

- A. Directory
- B. File
- C. Pre-signed
- D. Scoped

Answer: C

Explanation:

The pre-signed URL type allows users or applications to download or access files directly from a Snowflake stage without authentication. This URL type is open and can be used without needing to authenticate into Snowflake or pass an authorization token.

Question: 285

At what levels can a resource monitor be configured? (Select TWO).

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

Answer: A, E

Explanation:

Resource monitors in Snowflake can be configured at the account and virtual warehouse levels. They are used to track credit usage and control costs associated with running virtual warehouses. When certain thresholds are reached, resource monitors can trigger actions such as sending alerts or suspending warehouses to prevent excessive credit consumption. References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 286

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

Explanation:

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-C02] SnowPro Core Certification Exam Study Guide

Question: 287

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

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

Answer: A

Explanation:

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

Question: 288

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

Explanation:

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.

Question: 289

A tabular User-Defined Function (UDF) is defined by specifying a return clause that contains which keyword?

- A. ROW_NUMBER
- B. TABLE
- C. TABULAR
- D. VALUES

Answer: B

Explanation:

In Snowflake, a tabular User-Defined Function (UDF) is defined with a return clause that includes the keyword "TABLE." This indicates that the UDF will return a set of rows, which can be used in the FROM clause of a query. References: Based on my internal knowledge as of 2021.

Question: 290

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

- A. ALTER FILE FORMAT my_format SET TYPE=JSON;
- B. ALTER FILE FORMAT my_format SWAP TYPE WITH JSON;
- C. CREATE OR REPLACE FILE FORMAT my_format TYPE=JSON;
- D. REPLACE FILE FORMAT my_format TYPE=JSON;

Answer: A

Explanation:

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

Question: 291

What technique does Snowflake use to limit the number of micro-partitions scanned by each query?

- A. B-tree
- B. Indexing
- C. Map reduce
- D. Pruning

Answer: D

Explanation:

Snowflake uses a technique called pruning to limit the number of micro-partitions scanned by each query.

[Pruning effectively filters out unnecessary micro-partitions based on the query's filter conditions, which can significantly improve query performance by reducing the amount of data scanned1.](#)

Question: 292

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

Explanation:

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 use23.](#)

Question: 293

Which feature is integrated to support Multi-Factor Authentication (MFA) at Snowflake?

- A. Authy
- B. Duo Security
- C. One Login
- D. RSA SecurID Access

Answer: B

Explanation:

Snowflake integrates Duo Security to support Multi-Factor Authentication (MFA). [This feature provides increased login security for users connecting to Snowflake, and it is managed completely by Snowflake without the need for users to sign up separately with Duo4.](#)

Question: 294

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

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

Answer: A

Explanation:

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

Question: 295

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: C

Explanation:

The Snowflake object that helps evaluate virtual warehouse performance impacted by query queuing is the Information_schema.warehouse_load_history. [This view provides historical data about the load on a warehouse, including the average number of queries that were running or queued within a specific interval, which can be used to assess performance and identify potential issues with query queuing3.](#)

Question: 296

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

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

Answer: A

Explanation:

To find out which numeric values are present in both tables, the INTERSECT set operator should be used. [This operator returns rows from one query's result set which also appear in another query's result set, effectively finding the common elements between the two tables](#).

Question: 297

What type of columns does Snowflake recommend to be used as clustering keys? (Select TWO).

- A. A VARIANT column
- B. A column with very low cardinality
- C. A column with very high cardinality
- D. A column that is most actively used in selective filters
- E. A column that is most actively used in join predicates

Answer: C, D

Explanation:

Snowflake recommends using columns with very high cardinality and those that are most actively used in selective filters as clustering keys. High cardinality columns have a wide range of unique values, which helps in evenly distributing the data across micro-partitions. Columns used in selective filters help in pruning the number of micro-partitions to scan, thus improving query performance. References: Based on general database optimization principles.

Question: 298

Which of the following describes the Snowflake Cloud Services layer?

- A. Coordinates activities in the Snowflake account

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

Answer: A

Explanation:

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

Question: 299

What does Snowflake recommend regarding database object ownership? (Select TWO).

- A. Create objects with ACCOUNTADMIN and do not reassign ownership.
- B. Create objects with SYSADMIN.
- C. Create objects with SECURITYADMIN to ease granting of privileges later.
- D. Create objects with a custom role and grant this role to SYSADMIN.
- E. Use only MANAGED ACCESS SCHEMAS for objects owned by ACCOUNTADMIN.

Answer: B, D

Explanation:

Snowflake recommends creating objects with a role that has the necessary privileges and is not overly permissive. SYSADMIN is typically used for managing system-level objects and operations. Creating objects with a custom role and granting this role to SYSADMIN allows for more granular control and adherence to the principle of least privilege. References: Based on best practices for database object ownership and role management.

Question: 300

If a multi-cluster warehouse is using an economy scaling policy, how long will queries wait in the queue before another cluster is started?

- A. 1 minute
- B. 2 minutes
- C. 6 minutes

D. 8 minutes

Answer: B

Explanation:

In a multi-cluster warehouse with an economy scaling policy, queries will wait in the queue for 2 minutes before another cluster is started. This is to minimize costs by allowing queries to queue up for a short period before adding additional compute resources. References: [COF-C02] SnowPro Core Certification Exam

Study Guide

Question: 301

How can a Snowflake user access a JSON object, given the following table? (Select TWO).

SRC

```
(
  "salesperson": {
    "id": "1234",
    "name": "user"
  }
)
H
```

- A. src:salesperson.name
- B. src:salesPerson.name
- C. src:salesperson.Name
- D. SRC:salesperson.name
- E. SRC:salesperson.Name

Answer: A, C

Explanation:

To access a JSON object in Snowflake, dot notation is used where the path to the object is specified after the column name containing the JSON data. Both lowercase and uppercase can be used for attribute names, so both "name" and "Name" are valid. References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 302

What happens when a database is cloned?

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

D. It replicates all granted privileges on the corresponding child schema objects.

Answer: A

Explanation:

When a database is cloned in Snowflake, it does not retain any privileges that were granted on the source object. The clone will need to have privileges reassigned as necessary for users to access it. References: [COF-CO2] SnowPro Core Certification Exam Study Guide

Question: 303

Which native data types are used for storing semi-structured data in Snowflake? (Select TWO)

- A. NUMBER
- B. OBJECT
- C. STRING
- D. VARCHAR
- E. VARIANT

Answer: B, E

Explanation:

Snowflake supports semi-structured data types, which include OBJECT and VARIANT. These data types are capable of storing JSON-like data structures, allowing for flexibility in data representation. [OBJECT can directly contain VARIANT, and thus indirectly contain any other data type, including itself1.](#)

Question: 304

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

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

Answer: A, C

Explanation:

Snowflake's hierarchical key model includes several levels of keys, where Account master keys and File keys are part of this hierarchy. [Account master keys are used to encrypt all the data within an account, while File keys are used to encrypt individual files within the database2.](#)

Question: 305

How does Snowflake recommend handling the bulk loading of data batches from files already available in cloud storage?

- A. Use Snowpipe.
- B. Use the INSERT command.
- C. Use an external table.
- D. Use the COPY command.

Answer: D

Explanation:

Snowflake recommends using the COPY command for bulk loading data batches from files already available in cloud storage. [This command allows for efficient and large-scale data loading operations](#)

[from files staged in cloud storage into Snowflake tables3.](#)

Question: 306

What role is required to use Partner Connect?

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

Answer: A

Explanation:

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

Question: 307

How does a scoped URL expire?

- A. When the data cache clears.
- B. When the persisted query result period ends.
- C. The encoded URL access is permanent.
- D. The length of time is specified in the expiration_time argument.

Answer: B

Explanation:

A scoped URL expires when the persisted query result period ends, which is typically after the results cache expires. [This is currently set to 24 hours](#)

Question: 308

Which features make up Snowflake's column level security? (Select TWO).

- A. Continuous Data Protection (CDP)
- B. Dynamic Data Masking
- C. External Tokenization
- D. Key pair authentication
- E. Row access policies

Answer: B, C

Explanation:

Snowflake's column level security features include Dynamic Data Masking and External Tokenization. [Dynamic Data Masking uses masking policies to selectively mask data at query time, while External Tokenization allows for the tokenization of data before loading it into Snowflake and detokenizing it at query runtime.](#)

Question: 309

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

- A. Stored procedures can execute database operations while UDFs cannot.
- B. Returning a value is required in a stored procedure while returning values in a UDF is optional.
- C. Values returned by a stored procedure can be used directly in a SQL statement while the values returned

by a UDF cannot.

D. Multiple stored procedures can be called as part of a single executable statement while a single SQL statement can only call one UDF at a time.

Answer: A

Explanation:

[Stored procedures in Snowflake can perform a variety of database operations, including DDL and DML, whereas UDFs are designed to return values and cannot execute database operations1.](#)

Question: 310

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

Explanation:

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](#)

Question: 311

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

- A. It assigns each unloaded data file a unique name.
- B. It uses the put command to download the data by default.
- C. It uses COPY INTO <location> for bulk unloading where the default option is SINGLE - TRUE.
- D. It uses COPY INTO <location> to copy the data from a table into one or more files in an external stage only.

Answer: A

Explanation:

When unloading data, Snowflake assigns each file a unique name to ensure there is no overlap or confusion between files. [This is part of the bulk unloading process where data is exported from Snowflake tables into flat files3.](#)

Question: 312

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

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

Answer: D

Explanation:

[When file format options are specified in multiple locations, the load operation applies the options in the following order of precedence: first, the COPY INTO TABLE statement; second, the stage definition; and third, the table definition1](#)

Question: 313

Which statement accurately describes a characteristic of a materialized view?

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

Answer: B

Explanation:

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

Question: 314

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

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

Answer: B

Explanation:

The stream type that can be used for tracking the records in external tables is 'External'. [This type of stream is specifically designed to track changes in external tables](#)

Question: 315

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

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

Answer: D

Explanation:

In the Activity section in Snowsight, Snowflake users can explore the performance of executed queries. [This includes monitoring queries, viewing details about queries, including performance data, and exploring each step of an executed query in the query profile1.](#)

Question: 316

Which Snowflake URL type is used by directory tables?

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

Answer: C

Explanation:

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

Question: 317

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

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

Answer: D

Explanation:

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

Question: 318

What does Snowflake's search optimization service support?

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

Answer: C

Explanation:

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

Question: 319

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

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

Answer: C

Explanation:

To address query concurrency issues, a user can add additional clusters to the virtual warehouse. [This allows for the distribution of queries across multiple clusters, reducing the load on any single cluster and improving overall query performance.](#)

Question: 320

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: D

Explanation:

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.](#)

Question: 321

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: D, E

Explanation:

[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 data¹. It also includes the Automatic Clustering History, which provides insights into the automatic clustering operations performed on tables².](#)

Question: 322

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

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

Answer: A

Explanation:

The kind of Snowflake table that stores file-level metadata for each file in a stage is a directory table. [A directory table is an implicit object layered on a stage and stores file-level metadata about the data files in the stage³.](#)

Question: 323

Which parameter prevents streams on tables from becoming stale?

- A. MAXDATAEXTENSIONTIMEINDAYS
- B. MTN_DATA_RETENTION_TTIME_TN_DAYS
- C. LOCK_TIMEOUT
- D. STALE_AFTER

Answer: A

Explanation:

The parameter that prevents streams on tables from becoming stale is MAXDATAEXTENSIONTIMEINDAYS. [This parameter specifies the maximum number of days for which Snowflake can extend the data retention period for the table to prevent streams on the table from becoming stale⁴.](#)

Question: 324

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

A. GRANT ROLE MYROLE TO USER USER1;

GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO ROLE MYROLE;

B. GRANT ROLE MYROLE TO USER USER1;

GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER USER1;

C. GRANT ROLE MYROLE TO USER USER1;

GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB. K"-SCHEMA TO USER! ;

D. GRANT ROLE MYROLE TO USER USER1;

GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO MYROLE;

Answer: A

Explanation:

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

Question: 325

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

A. Account

B. Database

C. Schema

D. Table

E. Virtual warehouse

Answer: B, C

Explanation:

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

Question: 326

Two users share a virtual warehouse named wh dev 01. When one of the users loads data, the other One experiences performance issues while querying data.

How does Snowflake recommend resolving this issue?

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

Answer: C

Explanation:

Snowflake recommends creating separate warehouses for each workload to resolve performance issues caused by shared virtual warehouses. [This ensures that the resources are not being overutilized by one user's activities, thereby affecting the performance of another user's activities.](#)

Question: 327

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

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

Answer: D

Explanation:

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

Question: 328

A view is defined on a permanent table. A temporary table with the same name is created in the same schema as the referenced table. What will the query from the view return?

- A. The data from the permanent table.
- B. The data from the temporary table.

- C. An error stating that the view could not be compiled.
- D. An error stating that the referenced object could not be uniquely identified.

Answer: A

Explanation:

When a view is defined on a permanent table, and a temporary table with the same name is created in the same schema, the query from the view will return the data from the permanent table. [Temporary tables are session-specific and do not affect the data returned by views defined on permanent tables2.](#)

Question: 329

What is used to diagnose and troubleshoot network connections to Snowflake?

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

Answer: A

Explanation:

SnowCD (Snowflake Connectivity Diagnostic Tool) is used to diagnose and troubleshoot network connections to Snowflake. [It runs a series of connection checks to evaluate the network connection to Snowflake](#)

Question: 330

Which REST API can be used with unstructured data?

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

Answer: C

Explanation:

[The REST API used with unstructured data in Snowflake is GET /api/files/, which retrieves \(downloads\) a](#)

[data file from an internal or external stage](#).

Question: 331

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

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

Answer: D

Explanation:

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

Question: 332

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

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

Answer: B

Explanation:

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

Question: 333

By definition, a secure view is exposed only to users with what privilege?

- A. IMPORT SHARE
- B. OWNERSHIP
- C. REFERENCES

D. USAGE

Answer: B

Explanation:

A secure view in Snowflake is exposed only to users with the OWNERSHIP privilege. [This privilege ensures that only authorized users who own the view, or roles that include ownership, can access the secure view](#)

Question: 334

What is a characteristic of the Snowflake Query Profile?

- A. It can provide statistics on a maximum number of 100 queries per week.
- B. It provides a graphic representation of the main components of the query processing.
- C. It provides detailed statistics about which queries are using the greatest number of compute resources.
- D. It can be used by third-party software using the Query Profile API.

Answer: B

Explanation:

The Snowflake Query Profile provides a graphic representation of the main components of the query processing. [This visual aid helps users understand the execution details and performance characteristics of their queries4.](#)

Question: 335

What statistical information in a Query Profile indicates that the query is too large to fit in memory? (Select TWO).

- A. Bytes spilled to local cache.
- B. Bytes spilled to local storage.
- C. Bytes spilled to remote cache.
- D. Bytes spilled to remote storage.
- E. Bytes spilled to remote metastore.

Answer: A, B

Explanation:

In a Query Profile, the statistical information that indicates a query is too large to fit in memory includes bytes spilled to local cache and bytes spilled to local storage. [These metrics suggest that the working data set of the query exceeded the memory available on the warehouse nodes, causing intermediate results to be written to disk](#)

Question: 336

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

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

Answer: B

Explanation:

The maximum Time Travel retention period for a transient table in Snowflake is 1 day. [This is the default and maximum duration for which Snowflake maintains the historical data for transient tables, allowing users to query data as it appeared at any point within the past 24 hours2.](#)

Question: 337

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

- A. list @bronzestage;
- B. select * from directory(@bronzestage);
- C. select metadata\$filename from @bronzestage;
- D. select * from table(information_schema.stage_directory_file_registration_history(stage name=>'bronzestage1));

Answer: B

Explanation:

The query that contains a Snowflake hosted file URL in a directory table for a stage named bronzestage is select * from directory(@bronzestage). [This query retrieves a list of all files on the stage along with metadata, including the Snowflake file URL for each file3.](#)

Question: 338

Which Snowflake edition enables data sharing only through Snowflake Support?

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

Answer: A

Explanation:

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.](#)

Question: 339

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

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

Answer: A

Explanation:

Snowsight automatically detects if a target account is in a different region and enables cross-cloud auto-fulfillment when using a paid listing on the Snowflake Marketplace. [This feature allows Snowflake to manage the replication of data products to consumer regions as needed, without manual intervention1.](#)

Question: 340

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: B, E

Explanation:

User-Defined Function (UDF) handlers must be written inline for Javascript and SQL. [These languages allow the UDF logic to be included directly within the SQL statement that creates the UDF2.](#)

Question: 341

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

Explanation:

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.](#)

Question: 342

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: A, B

Explanation:

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](#)

Question: 343

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: C, D

Explanation:

Snowpark is designed to bring the data programmability to Snowflake, enabling developers to write code in familiar languages like Scala, Java, and Python. It allows for the execution of these codes directly within Snowflake's virtual warehouses, eliminating the need for a separate cluster. [Additionally, Snowpark's compatibility with Spark allows users to leverage their existing Spark code with minimal changes¹.](#)

Question: 344

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

- A. 60

B. 91

C. 120

D. 121

Answer: D

Explanation:

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

Question: 345

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

A. Filter using a where clause

B. Filter using a limit keyword

C. Cast using a SELECT statement

D. Order using an ORDER BY clause

Answer: C

Explanation:

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

Question: 346

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

Explanation:

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](#)

Question: 347

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

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

Answer: B

Explanation:

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

Topic 4, Exam pool D

Question: 348

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

- A. A user with an USERADMIN or higher role
- B. A user with a SECURITYADMIN or higher role
- C. A role that has been granted the ATTACH POLICY privilege
- D. A role that has the NETWORK_POLICY account parameter set
- E. A role that has the ownership of the network policy

Answer: B, E

Explanation:

In Snowflake, a user with the SECURITYADMIN role or higher can activate and enforce a network policy for all users in an account. [Additionally, a role that has ownership of the network policy can also activate and enforce it](#)

Question: 349

A permanent table and temporary table have the same name, TBL1, in a schema.

What will happen if a user executes select * from TBL1 ;?

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

Answer: A

Explanation:

[In Snowflake, if a temporary table and a permanent table have the same name within the same schema, the temporary table takes precedence over the permanent table within the session where the temporary table was created.](#)

Question: 350

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

Explanation:

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 semistructured data in Snowflake](#)

Question: 351

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

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

Answer: A

Explanation:

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

Question: 352

What is the purpose of a Query Profile?

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

Answer: B

Explanation:

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

Question: 353

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

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

Answer: B

Explanation:

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

Question: 354

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

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

Answer: C

Explanation:

The minimum Snowflake Edition that supports secure storage of Protected Health Information (PHI) data is the Business Critical Edition. [This edition offers enhanced security features necessary for compliance with regulations such as HIPAA and HITRUST CSF4.](#)

Question: 355

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

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

Answer: A, B

Explanation:

[Snowflake stores metadata for rows in micro-partitions, including the range of values for each column and the number of distinct values¹.](#)

Question: 356

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

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

Answer: C, D

Explanation:

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

Question: 357

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

Explanation:

[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.](#)

Question: 358

At what level is the MIN_DATA_RETENTION_TIME_IN_DAYS parameter set?

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

Answer: A

Explanation:

The MIN_DATA_RETENTION_TIME_IN_DAYS parameter is set at the account level. [This parameter determines the minimum number of days Snowflake retains historical data for Time Travel operations](#)

Question: 359

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

Explanation:

An account administrator in the Data Exchange can perform tasks such as adding and removing members and approving or denying listing approval requests. [These tasks are part of managing the Data Exchange and ensuring that only authorized listings and members are part of it12.](#)

Question: 360

Which statistics can be used to identify queries that have inefficient pruning? (Select TWO).

- A. Bytes scanned
- B. Bytes written to result
- C. Partitions scanned
- D. Partitions total
- E. Percentage scanned from cache

Answer: C, D

Explanation:

The statistics that can be used to identify queries with inefficient pruning are 'Partitions scanned' and 'Partitions total'. [These statistics indicate how much of the data was actually needed and scanned versus the total available, which can highlight inefficiencies in data pruning34.](#)

Question: 361

What is a directory table in Snowflake?

- A. A separate database object that is used to store file-level metadata
- B. An object layered on a stage that is used to store file-level metadata
- C. A database object with grantable privileges for unstructured data tasks
- D. A Snowflake table specifically designed for storing unstructured files

Answer: B

Explanation:

A directory table in Snowflake is an object layered on a stage that is used to store file-level metadata. [It is not a separate database object but is conceptually similar to an external table because it stores metadata about the data files in the stage5.](#)

Question: 362

Which Snowflake command can be used to unload the result of a query to a single file?

- A. Use COPY INTO <external stage> followed by a GET command to download the file.
- B. Use COPY INTO <internal stage> followed by a put command to download the file.
- C. Use COPY INTO <internal stage> with SINGLE = TRUE followed by a GET command to download the file.
- D. Use COPY INTO <external stage> with SINGLE = TRUE followed by a PUT command to download the file.

Answer: C

Explanation:

The Snowflake command to unload the result of a query to a single file is COPY INTO <internal stage> with SINGLE = TRUE followed by a GET command to download the file. [This command unloads the query result into a single file in the specified internal stage](#)

Question: 363

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

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

Answer: C

Explanation:

In a managed access schema, the schema owner retains the OWNERSHIP privilege of any tables added to the schema. [This means that while object owners have certain privileges over the objects they create, only the schema owner can manage privilege grants on these objects1.](#)

Question: 364

When using the ALLOW_CLIENT_MFA_CACHING parameter, how long is a cached Multi-Factor Authentication (MFA) token valid for?

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

Answer: C

Explanation:

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 timeframe2.](#)

Question: 365

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

Explanation:

The factors that impact storage costs in Snowflake include the account type (Capacity or On Demand) and the cloud region used by the account. [These factors determine the rate at which storage is billed, with different regions potentially having different rates3.](#)

Question: 366

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

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

Answer: D

Explanation:

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

Question: 367

Which ACCOUNT_USAGE schema database role provides visibility into policy-related information?

- A. USAGE_VIEWER
- B. GOVERNANCE_VIEWER
- C. OBJECT_VIEWER
- D. SECURITY_VIEWER

Answer: B

Explanation:

The GOVERNANCE_VIEWER role in the ACCOUNT_USAGE schema provides visibility into policy-related information within Snowflake. [This role is specifically designed to access views that display object metadata and usage metrics related to governance12.](#)

Question: 368

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

- A. INTEGER
- B. OBJECT
- C. VARIANT

D. VARCHAR

E. LIST

Answer: B, C

Explanation:

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

Question: 369

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

Explanation:

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](#)

Question: 370

Who can grant object privileges in a regular schema?

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

Answer: A

Explanation:

In a regular schema within Snowflake, the object owner has the privilege to grant object privileges. [The object owner is typically the role that created the object or to whom the ownership of the object has been transferred](#)⁷⁸.

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

Question: 371

Which command is used to unload data from a Snowflake database table into one or more files in a Snowflake stage?

- A. CREATE STAGE
- B. COPY INTO <table>
- C. COPY INTO <location>
- D. CREATE PIPE

Answer: C

Explanation:

[The COPY INTO <location> command is used to unload data from a Snowflake database table into one or more files in a Snowflake stage](#)¹.

Question: 372

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

- A. DATABASE_STORAGE_USAGE_HISTORY
- B. TABLE_STORAGE_METRICS
- C. STORAGE_DAILY_HISTORY
- D. STAGE_STORAGE_USAGE_HISTORY

Answer: B

Explanation:

[The TABLE_STORAGE_METRICS metadata table stores the storage utilization information, including for tables that have been dropped but are still incurring storage costs](#)².

Question: 373

What does a masking policy consist of in Snowflake?

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

Answer: A

Explanation:

A masking policy in Snowflake consists of a single data type, with one or more conditions, and one or more masking functions. [These components define how the data is masked based on the specified conditions](#)3.

Question: 374

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

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

Answer: D

Explanation:

[Snowflake's Time Travel feature is used for the maintenance of historical data, allowing users to access and restore data that has been changed or deleted within a defined period](#)4.

Question: 375

A JSON file, that contains lots of dates and arrays, needs to be processed in Snowflake. The user wants to ensure optimal performance while querying the data.

How can this be achieved?

- A. Flatten the data and store it in structured data types in a flattened table. Query the table.
- B. Store the data in a table with a variant data type. Query the table.

- C. Store the data in a table with a variant data type and include STRIP_NULL_VALUES while loading the table. Query the table.
- D. Store the data in an external stage and create views on top of it. Query the views.

Answer: B

Explanation:

Storing JSON data in a table with a VARIANT data type is optimal for querying because it allows Snowflake to leverage its semi-structured data capabilities. [This approach enables efficient storage and querying without the need for flattening the data, which can be performance-intensive1.](#)

Question: 376

For which use cases is running a virtual warehouse required? (Select TWO).

- A. When creating a table
- B. When loading data into a table
- C. When unloading data from a table
- D. When executing a show command
- E. When executing a list command

Answer: B, C

Explanation:

[Running a virtual warehouse is required when loading data into a table and when unloading data from a table because these operations require compute resources that are provided by the virtual warehouse23.](#)

Question: 377

What does SnowCD help Snowflake users to do?

- A. Copy data into files.
- B. Manage different databases and schemas.
- C. Troubleshoot network connections to Snowflake.
- D. Write SELECT queries to retrieve data from external tables.

Answer: C

Explanation:

SnowCD is a connectivity diagnostic tool that helps users troubleshoot network connections to Snowflake. [It performs a series of checks to evaluate the network connection and provides suggestions for resolving any issues4.](#)

Question: 378

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

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

Answer: D, E

Explanation:

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

Question: 379

How is unstructured data retrieved from data storage?

- A. SQL functions like the GET command can be used to copy the unstructured data to a location on the client.
- B. SQL functions can be used to create different types of URLs pointing to the unstructured data. These URLs can be used to download the data to a client.
- C. SQL functions can be used to retrieve the data from the query results cache. When the query results are output to a client, the unstructured data will be output to the client as files.
- D. SQL functions can call on different web extensions designed to display different types of files as a web page. The web extensions will allow the files to be downloaded to the client.

Answer: B

Explanation:

Unstructured data stored in Snowflake can be retrieved by using SQL functions to generate URLs that point to the data. [These URLs can then be used to download the data directly to a client](#)

Question: 380

Which Snowflake view is used to support compliance auditing?

- A. ACCESS_HISTORY
- B. COPY_HISTORY
- C. QUERY_HISTORY
- D. ROW ACCESS POLICIES

Answer: A

Explanation:

The ACCESS_HISTORY view in Snowflake is utilized to support compliance auditing. It provides detailed information on data access within Snowflake, including reads and writes by user queries. [This view is essential for regulatory compliance auditing as it offers insights into the usage of tables and columns, and maintains a direct link between the user, the query, and the accessed data1.](#)

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

Question: 381

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

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

Answer: D

Explanation:

The STRIP NULL_VALUES file format option, when set to TRUE, removes object or array elements that contain null values during the loading process of semi-structured data files into Snowflake. [This ensures that the data loaded into Snowflake tables does not contain these null elements, which can be useful when the “null” values in files indicate missing values and have no other special meaning.](#)

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

Question: 382

Which data types can be used in Snowflake to store semi-structured data? (Select TWO)

- A. ARRAY
- B. BLOB
- C. CLOB
- D. JSON
- E. VARIANT

Answer: A, E

Explanation:

Snowflake supports the storage of semi-structured data using the ARRAY and VARIANT data types. The ARRAY data type can directly contain VARIANT, and thus indirectly contain any other data type, including itself.

[The VARIANT data type can store a value of any other type, including OBJECT and ARRAY, and is often used to represent semi-structured data formats like JSON, Avro, ORC, Parquet, or XML.](#)

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

Question: 383

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

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

Answer: D

Explanation:

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

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

Question: 384

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

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

Answer: C

Explanation:

The Notify & Suspend action for a resource monitor in Snowflake sends a notification to all account administrators who have notifications enabled and suspends all assigned warehouses. [However, the suspension only occurs after all currently running statements in the warehouses have been completed1.](#)

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

Question: 385

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

Explanation:

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](#)². References: [COF- C02] SnowPro Core Certification Exam Study Guide

Question: 386

What Snowflake feature provides a data hub for secure data collaboration, with a selected group of invited members?

- A. Data Replication
- B. Secure Data Sharing
- C. Data Exchange
- D. Snowflake Marketplace

Answer: C

Explanation:

Snowflake's Data Exchange feature provides a data hub for secure data collaboration. [It allows providers to publish data that can be discovered and accessed by a selected group of invited members, facilitating secure and controlled data sharing within a collaborative environment](#)³. References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 387

What is a characteristic of materialized views in Snowflake?

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

Answer: C

Explanation:

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

Question: 388

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

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

Answer: C

Explanation:

To start configuring Snowflake for Single Sign-On (SSO), the CREATE SECURITY INTEGRATION command is used. [This command sets up a security integration object in Snowflake, which is necessary for enabling SSO with external identity providers using SAML 2.01](#).

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

Question: 389

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

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

Answer: A

Explanation:

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

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

Question: 390

Which operation can be performed on Snowflake external tables?

- A. INSERT
- B. JOIN
- C. RENAME
- D. ALTER

Answer: B

Explanation:

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 operations3.](#)

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

Question: 391

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

Explanation:

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

Question: 392

What step can reduce data spilling in Snowflake?

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

Answer: A

Explanation:

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

Question: 393

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

Explanation:

[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](#)². References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 394

Which solution improves the performance of point lookup queries that return a small number of rows from large tables using highly selective filters?

- A. Automatic clustering
- B. Materialized views
- C. Query acceleration service
- D. Search optimization service

Answer: D

Explanation:

The search optimization service improves the performance of point lookup queries on large tables by using selective filters to quickly return a small number of rows. [It creates an optimized data structure that helps in pruning the micro-partitions that do not contain the queried values](#)³. References: [COF- C02] SnowPro Core Certification Exam Study Guide

Question: 395

Which Snowflake object does not consume any storage costs?

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

Answer: C

Explanation:

[Temporary tables do not consume any storage costs in Snowflake because they only exist for the duration of the session that created them and are automatically dropped when the session ends, thus incurring](#)

[no long-term storage charges4](#). References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 396

Which function unloads data from a relational table to JSON?

- A. TO_OBJECT
- B. TO_JSON
- C. TO_VARIANT
- D. OBJECT CONSTRUCT

Answer: B

Explanation:

The TO_JSON function is used to convert a VARIANT value into a string containing the JSON representation of the value. This function is suitable for unloading data from a relational table to JSON format. References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 397

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

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

Answer: C, D

Explanation:

The LIST and PUT commands are specific to SnowSQL and cannot be executed in the web interface or other SQL clients. LIST is used to display the contents of a stage, and PUT is used to upload files to a

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

Question: 398

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

Explanation:

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-C02] SnowPro Core Certification Exam Study Guide

Question: 399

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

Explanation:

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

Question: 400

Which statements describe benefits of Snowflake's separation of compute and storage? (Select TWO).

- A. The separation allows independent scaling of computing resources.

- B. The separation ensures consistent data encryption across all virtual data warehouses.
- C. The separation supports automatic conversion of semi-structured data into structured data for advanced data analysis.
- D. Storage volume growth and compute usage growth can be tightly coupled.
- E. Compute can be scaled up or down without the requirement to add more storage.

Answer: A, E

Explanation:

Snowflake's architecture allows for the independent scaling of compute resources, meaning you can increase or decrease the computational power as needed without affecting storage. This separation also means that storage can grow independently of compute usage, allowing for more flexible and cost-effective data management.

Question: 401

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

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

Answer: A, E

Explanation:

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

Question: 402

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

- A. Access history

- B. Data classification
- C. Dynamic Data Masking
- D. Object tagging

Answer: A

Explanation:

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

Question: 403

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

Explanation:

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.

Question: 404

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

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

- D. Use the query acceleration service

Answer: C

Explanation:

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

Question: 405

What does the LATERAL modifier for the FLATTEN function do?

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

Answer: C

Explanation:

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

Question: 406

Which Snowflake feature allows a user to track sensitive data for compliance, discovery, protection, and resource usage?

- A. Tags
- B. Comments
- C. Internal tokenization
- D. Row access policies

Answer: A

Explanation:

Tags in Snowflake allow users to track sensitive data for compliance, discovery, protection, and resource usage.

[They enable the categorization and tracking of data, supporting compliance with privacy regulations678.](#)

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

Question: 407

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: B, C

Explanation:

Virtual warehouses in Snowflake can be started and stopped at any time, providing flexibility in managing compute resources. [They can also be resized at any time, even while running, to accommodate varying workloads910.](#) References: [COF-C02] SnowPro Core Certification Exam Study Guide

Question: 408

Which parameter can be set at the account level to set the minimum number of days for which

Snowflake retains historical data in Time Travel?

- A. DATA_RETENTION_TIME_IN_DAYS
- B. MAX_DATA_EXTENSION_TIME_IN_DAYS
- C. MIN_DATA_RETENTION_TIME_IN_DAYS
- D. MAX_CONCURRENCY_LEVEL

Answer: A

Explanation:

[The parameter DATA_RETENTION_TIME_IN_DAYS can be set at the account level to define the minimum number of days Snowflake retains historical data for Time Travel1.](#)

Question: 409

A tag object has been assigned to a table (TABLE_A) in a schema within a Snowflake database.

Which CREATE object statement will automatically assign the TABLE_A tag to a target object?

- A. CREATE TABLE <table_name> LIKE TABLE_A;
- B. CREATE VIEW <view_name> AS SELECT * FROM TABLE_A;
- C. CREATE TABLE <table_name> AS SELECT * FROM TABLE_A;
- D. CREATE MATERIALIZED VIEW <view name> AS SELECT * FROM TABLE A;

Answer: C

Explanation:

[When a tag object is assigned to a table, using the statement CREATE TABLE <table_name> AS SELECT * FROM TABLE_A will automatically assign the TABLE_A tag to the newly created table2.](#)

Question: 410

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

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

Answer: A, C

Explanation:

Dynamic Data Masking in Snowflake supports tables and views. [These objects can have masking policies](#)

[applied to their columns to dynamically mask data at query time](#)³.

Question: 411

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

Explanation:

[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](#)⁴.

Question: 412

Which Snowflake function will parse a JSON-null into a SQL-null?

- A. TO_CHAR
- B. TO_VARIANT
- C. TO_VARCHAR
- D. STRIP NULL VALUE

Answer: D

Explanation:

[The STRIP NULL VALUE function in Snowflake is used to convert a JSON null value into a SQL NULL value](#)¹.

Question: 413

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

- A. GET_PRESIGNED_URL

- B. GET_ABSOLUTE_PATH
- C. BUILD_STAGE_FILE_URL
- D. BUILD SCOPED FILE URL

Answer: C

Explanation:

[The BUILD_STAGE_FILE_URL function generates a Snowflake-hosted file URL to a staged file using the stage name and relative file path as inputs2.](#)

Question: 414

Which Snowflake database object can be used to track data changes made to table data?

- A. Tag
- B. Task
- C. Stream
- D. Stored procedure

Answer: C

Explanation:

[A Stream object in Snowflake is used for change data capture \(CDC\), which records data manipulation language \(DML\) changes made to tables, including inserts, updates, and deletes3.](#)

Question: 415

What transformations are supported when loading data into a table using the COPY INTO <table> command? (Select TWO).

- A. Column reordering
- B. Column omission
- C. JOIN function
- D. FLATTEN function

E. GROUP BY function

Answer: A, B

Explanation:

[The COPY INTO <table> command in Snowflake supports column reordering and column omission as part of its data transformation capabilities during the load process⁴⁵.](#)

Question: 416

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

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

Answer: B, D

Explanation:

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

Question: 417

Which object can be used with Secure Data Sharing?

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

Answer: A

Explanation:

Views can be used with Secure Data Sharing in Snowflake. [Materialized views, external tables, and UDFs are](#)

[not typically shared directly for security and performance reasons2.](#)

Question: 418

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

- A. Review the granted privileges to the database object.
- B. Review the row access policy for the database object.
- C. Query the ACCESS_HISTORY view in the ACCOUNT_USAGE schema.
- D. Query the REPLICATION USAGE HISTORY view in the ORGANIZATION USAGE schema.

Answer: C

Explanation:

[To determine which user has accessed a database object containing sensitive information, a Snowflake administrator can query the ACCESS_HISTORY view in the ACCOUNT_USAGE schema, which provides information about access to database objects3.](#)

Question: 419

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

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

Answer: B

Explanation:

[The LOOP type of loop in Snowflake Scripting does not have a built-in termination condition and requires a BREAK statement to stop executing4.](#)

Question: 420

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: C

Explanation:

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-partitions1.](#)

Question: 421

Which views are included in the DATA SHARING USAGE schema? (Select TWO).

- A. ACCESS_HISTORY
- B. DATA_TRANSFER_HISTORY
- C. WAREHOUSE_METERING_HISTORY
- D. MONETIZED_USAGE_DAILY
- E. LISTING_TELEMETRY_DAILY

Answer: D, E

Explanation:

[The DATA_SHARING_USAGE schema includes views that display information about listings published in the Snowflake Marketplace or a data exchange, which includes DATA_TRANSFER_HISTORY and LISTING_TELEMETRY_DAILY2.](#)

Question: 422

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

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

Answer: C, D

Explanation:

To view and modify resource monitors, the least privileges needed are MONITOR and MODIFY. [These privileges allow a user to monitor credit usage and make changes to resource monitors3.](#)

Question: 423

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

Explanation:

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 column4.](#)

Question: 424

What is the purpose of the Snowflake SPLIT TO_TABLE function?

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

Answer: C

Explanation:

The purpose of the Snowflake SPLIT_TO_TABLE function is to split a string based on a specified delimiter and flatten the results into rows. [This table function is useful for transforming a delimited string into a set of rows that can be further processed or queried5.](#)

Question: 425

A column named "Data" contains VARIANT data and stores values as follows:

```
"Employee" : { id : 100, name : John, location : "New York"
}
```

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: D

Explanation:

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.

Question: 426

Which Snowflake feature provides increased login security for users connecting to Snowflake that is powered by Duo Security service?

- A. OAuth
- B. Network policies

- C. Single Sign-On (SSO)
- D. Multi-Factor Authentication (MFA)

Answer: D

Explanation:

Multi-Factor Authentication (MFA) provides increased login security for users connecting to Snowflake. Snowflake's MFA is powered by Duo Security service, which adds an additional layer of security during the login process.

Question: 427

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

Explanation:

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.

Question: 428

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: D

Explanation:

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.

Question: 429

What information is found within the Statistic output in the Query Profile Overview?

A. Operator tree

B. Table pruning

C. Most expensive nodes

D. Nodes by execution time

Answer: C

Explanation:

The Statistic output in the Query Profile Overview of Snowflake provides detailed insights into the performance of different parts of the query. Specifically, it highlights the "Most expensive nodes," which are the operations or steps within the query execution that consume the most resources, such as CPU and memory. Identifying these nodes helps in pinpointing performance bottlenecks and optimizing query execution by focusing efforts on the most resource-intensive parts of the query.

References:

Snowflake Documentation on Query Profile Overview: It details the components of the profile overview, emphasizing how to interpret the statistics section to improve query performance by understanding which

nodes are most resource-intensive.

Question: 430

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: D

Explanation:

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 nonsecure views, clarifying that the main difference lies in security features rather than performance, thus supporting the assertion that there are no inherent performance differences.

Question: 431

When using SnowSQL, which configuration options are required when unloading data from a SQL query run on a local machine? (Select TWO).

- A. echo
- B. quiet
- C. output_file
- D. output_format
- E. force_put_overwrite

Answer: C, D

Explanation:

When unloading data from SnowSQL (Snowflake's command-line client), to a file on a local machine, you need to specify certain configuration options to determine how and where the data should be outputted. The correct configuration options required are:

C . `output_file`: This configuration option specifies the file path where the output from the query should be stored. It is essential for directing the results of your SQL query into a local file, rather than just displaying it on the screen.

D . `output_format`: This option determines the format of the output file (e.g., CSV, JSON, etc.). It is crucial for ensuring that the data is unloaded in a structured format that meets the requirements of downstream processes or systems.

These options are specified in the SnowSQL configuration file or directly in the SnowSQL command line. The configuration file allows users to set defaults and customize their usage of SnowSQL, including output preferences for unloading data.

References:

Snowflake Documentation: SnowSQL (CLI Client) at Snowflake Documentation

Snowflake Documentation: Configuring SnowSQL at Snowflake Documentation

Question: 432

How can a Snowflake user post-process the result of SHOW FILE FORMATS?

- A. Use the `RESULT_SCAN` function.
- B. Create a `CURSOR` for the command.
- C. Put it in the `FROM` clause in brackets.
- D. Assign the command to `RESULTSET`.

Answer: A

Explanation:

first run `SHOW FILE FORMATS`

then `SELECT * FROM TABLE(RESULT_SCAN(LAST_QUERY_ID(-1)))`

https://docs.snowflake.com/en/sql-reference/functions/result_scan#usage-notes

Question: 433

Which file function gives a user or application access to download unstructured data from a Snowflake stage?

- A. BUILD_SCOPED_FILE_URL
- B. BUILD_STAGE_FILE_URL
- C. GET_PRESIGNED_URL
- D. GET_STAGE_LOCATION

Answer: C

Explanation:

The function that provides access to download unstructured data from a Snowflake stage is:

C. GET_PRESIGNED_URL: This function generates a presigned URL for a single file within a stage. The generated URL can be used to directly access or download the file without needing to go through Snowflake. This is particularly useful for unstructured data such as images, videos, or large text files, where direct access via a URL is needed outside of the Snowflake environment.

Example usage:

```
SELECT GET_PRESIGNED_URL('stage_name', 'file_path');
```

This function simplifies the process of securely sharing or accessing files stored in Snowflake stages with external systems or users.

References:

Snowflake Documentation: GET_PRESIGNED_URL Function at Snowflake Documentation

Question: 434

When should a multi-cluster virtual warehouse be used in Snowflake?

- A. When queuing is delaying query execution on the warehouse
- B. When there is significant disk spilling shown on the Query Profile
- C. When dynamic vertical scaling is being used in the warehouse
- D. When there are no concurrent queries running on the warehouse

Answer: A

Explanation:

A multi-cluster virtual warehouse in Snowflake is designed to handle high concurrency and workload demands by allowing multiple clusters of compute resources to operate simultaneously. The correct scenario to use a multi-cluster virtual warehouse is:

A . When queuing is delaying query execution on the warehouse: Multi-cluster warehouses are ideal when the demand for compute resources exceeds the capacity of a single cluster, leading to query queuing. By enabling additional clusters, you can distribute the workload across multiple compute clusters, thereby reducing queuing and improving query performance.

This is especially useful in scenarios with fluctuating workloads or where it's critical to maintain low response times for a large number of concurrent queries.

References:

Question: 435

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

Explanation:

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.

Question: 436

Which Snowflake tool is recommended for data batch processing?

- A. SnowCD
- B. SnowSQL
- C. Snowsight
- D. The Snowflake API

Answer: B

Explanation:

For data batch processing in Snowflake, the recommended tool is:

B. SnowSQL: SnowSQL is the command-line client for Snowflake. It allows for executing SQL queries, scripts, and managing database objects. It's particularly suitable for batch processing tasks due to its ability to run SQL scripts that can execute multiple commands or queries in sequence, making it ideal for automated or scheduled tasks that require bulk data operations.

SnowSQL provides a flexible and powerful way to interact with Snowflake, supporting operations such as loading and unloading data, executing complex queries, and managing Snowflake objects from the command line or through scripts.

References:

Snowflake Documentation: SnowSQL (CLI Client) at Snowflake Documentation

Question: 437

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

Explanation:

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

Question: 438

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

Explanation:

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

Question: 439

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: C

Explanation:

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

Question: 440

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

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

Answer: D

Explanation:

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

Question: 441

In Snowflake, the use of federated authentication enables which Single Sign-On (SSO) workflow activities? (Select TWO).

- A. Authorizing users
- B. Initiating user sessions
- C. Logging into Snowflake
- D. Logging out of Snowflake
- E. Performing role authentication

Answer: B C

Explanation:

Federated authentication in Snowflake allows users to use their organizational credentials to log in to Snowflake, leveraging Single Sign-On (SSO). The key activities enabled by this setup include:

- B . Initiating user sessions: Federated authentication streamlines the process of starting a user session in Snowflake by using the existing authentication mechanisms of an organization.
- C . Logging into Snowflake: It simplifies the login process, allowing users to authenticate with their organization's identity provider instead of managing separate credentials for Snowflake.

References:

Snowflake Documentation: Configuring Federated Authentication

Question: 442

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: B, E

Explanation:

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

Topic 5, Exam pool E

Question: 443

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 (
EMP_ID NUMBER,
EMP_NAME VARCHAR(30),
EMP_SALARY NUMBER, DEPT VARCHAR(20)
);`

Answer: C

Explanation:

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

Question: 444

Which statement accurately describes how a virtual warehouse functions?

- A. Increasing the size of a virtual warehouse will always improve data loading performance.
- B. Each virtual warehouse is an independent compute cluster that shares compute resources with other warehouses.
- C. Each virtual warehouse is a compute cluster composed of multiple compute nodes allocated by Snowflake from a cloud provider.
- D. All virtual warehouses share the same compute resources so performance degradation of one warehouse can significantly affect all the other warehouses.

Answer: C

Explanation:

A virtual warehouse in Snowflake is an independent compute cluster that performs data processing tasks such as executing SQL queries. Each virtual warehouse is dynamically allocated by Snowflake from the cloud provider's resources and does not share compute resources with other warehouses. This architecture ensures that the performance of one warehouse does not impact the performance of another. Adjusting the size of a virtual warehouse affects its computational power by increasing or decreasing the number of compute nodes, which can improve the performance of data processing tasks depending on the workload.

References:

Snowflake Documentation: Understanding Virtual Warehouses

Question: 445

From what stage can a Snowflake user omit the FROM clause while loading data into a table?

- A. The user stage
- B. The table stage
- C. The internal named stage
- D. The external named stage

Answer: B

Explanation:

In Snowflake, when loading data into a table using the COPY INTO command, the FROM clause can be omitted if loading from the table's stage, also known as the table stage. The table stage is a default location associated with each table where files can be temporarily stored for loading operations. This simplifies the data loading process by allowing direct loading from files that have been uploaded to the table's stage without specifying the stage explicitly in the COPY INTO

command.

References:

Snowflake Documentation: Loading Data into Tables

Question: 446

A Snowflake user needs to share unstructured data from an internal stage to a reporting tool that does not have

Snowflake access.

Which file function should be used?

- A. BUILD_SCOPED_FILE_URL
- B. BUILD_STAGE_FILE_URL
- C. GET_PRESIGNED_URL
- D. GET_STAGE_LOCATION

Answer: C

Explanation:

The GET_PRESIGNED_URL function in Snowflake generates a presigned URL for a file stored in an internal stage, allowing direct access to the file without requiring Snowflake access. This feature is particularly useful for sharing unstructured data with external applications or tools that do not have direct access to Snowflake. The presigned URL provides temporary access to the file, making it an ideal solution for securely sharing unstructured data from an internal stage with a reporting tool or any other external application.

References:

Snowflake Documentation: Generating Presigned URLs

Question: 447

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

Explanation:

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

Question: 448

What is a feature of column-level security in Snowflake?

- A. Role access policies
- B. Network policies
- C. Internal tokenization
- D. External tokenization

Answer: A

Explanation:

Column-level security in Snowflake is implemented through Role Access Policies. These policies allow administrators to control access to specific columns of a table or view based on the role of the user accessing the data. By applying a role access policy to a column, administrators can ensure that sensitive information remains secure, and only users with the appropriate roles can view or query the data in that column. This feature enhances the security model by providing fine-grained access control at the column level.

References:

Question: 449

What operations can be performed while loading a simple CSV file into a Snowflake table using the COPY INTO command? (Select TWO).

- A. Performing aggregate calculations
- B. Reordering the columns
- C. Grouping by operations
- D. Converting the datatypes
- E. Selecting the first few rows

Answer: B, D

Explanation:

When loading a simple CSV file into a Snowflake table using the COPY INTO command, you can perform various transformations and adjustments on the data as part of the loading process. Specifically, you can:

B . Reorder the columns: Specify the order of columns in the COPY INTO command to match the structure of the target table if it differs from the order of columns in the source CSV file.

D . Convert the datatypes: Explicitly convert the datatypes of the data being loaded to match the datatypes of the columns in the target table. This can be necessary when the source data's format does not match the target table's expected datatype.

References:

Snowflake Documentation: Using the COPY INTO Command for Data Loading

Question: 450

What consideration should be made when loading data into Snowflake?

- A. Create small data files and stage them in cloud storage frequently.
- B. Create large data files to maximize the processing overhead for each file.
- C. The number of load operations That run in parallel can exceed the number of data files to be loaded.
- D. The number of data files that are processed in parallel is determined by the virtual warehouse.

Answer: D

Explanation:

When loading data into Snowflake, one critical consideration is the parallel processing capability of the virtual warehouse used for the data loading operation. The number of data files that can be processed in parallel during a loading operation is determined by the size and resources of the virtual warehouse. A larger warehouse can process more files simultaneously, improving the efficiency and speed of data loading operations. Optimizing the size of the virtual warehouse according to the data loading needs and the size and number of files to be loaded can significantly impact the overall performance of the data loading process.

References:

Snowflake Documentation: Optimizing Data Loading

Question: 451

Which Snowflake feature can be used to find sensitive data in a table or column?

- A. Masking policies
- B. Data classification
- C. Row level policies
- D. External functions

Answer: B

Explanation:

Data classification in Snowflake is a feature that allows organizations to identify and categorize data stored in tables or columns based on its sensitivity level or content type. This feature can be used to find sensitive data within the database by classifying data as confidential, personal, public, etc., making it easier to apply appropriate security measures, such as masking policies or row-level security, to protect sensitive information.

References:

Snowflake Documentation: Data Classification

Question: 452

What can the Snowflake SCIM API be used to manage? (Select TWO).

- A. Integrations
- B. Network policies
- C. Session policies

D. Roles

E. Users

Answer: D, E

Explanation:

The Snowflake SCIM (System for Cross-domain Identity Management) API is used for automated user and role management. It enables integration with identity providers (IdPs) for the provisioning and deprovisioning of user accounts and roles in Snowflake. This helps in managing access control and permissions systematically and aligns with identity governance practices.

References:

Snowflake Documentation: Managing Users and Roles with SCIM API

Question: 453

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

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

Answer: D

Explanation:

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

References:

Snowflake Documentation: Querying Semi-structured Data

Question: 454

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

Explanation:

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

their permissions, enhancing security and governance within Snowflake environments.

References:

Snowflake Documentation: Understanding Roles

Question: 455

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

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

Answer: C

Explanation:

In Snowflake, overloading refers to the creation of multiple User-Defined Functions (UDFs) with the same name but differing in the number or types of their arguments. This feature allows for more flexible function usage, as Snowflake

can differentiate between functions based on the context of their invocation, such as the types or the number of arguments passed. Overloading helps to create more adaptable and readable code, as the same function name can be used for similar operations on different types of data.

References:

Snowflake Documentation: User-Defined Functions

Question: 456

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

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

Answer: A, E

Explanation:

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

References:

Snowflake Documentation: Semi-structured Data Types

Question: 457

Which SQL command can be used to verify the privileges that are granted to a role?

- A. SHOW GRANTS ON ROLE <Role Name>
- B. SHOW ROLES <Role Name>
- C. SHOW GRANTS TO ROLE <Role Name>
- D. SHOW GRANTS FOR ROLE <Role Name>

Answer: C

Explanation:

To verify the privileges that have been granted to a specific role in Snowflake, the correct SQL command is `SHOW GRANTS TO ROLE <Role Name>`. This command lists all the privileges granted to the specified role, including access to schemas, tables, and other database objects. This is a useful command for administrators and users with sufficient privileges to audit and manage role permissions within the Snowflake environment.

References:

Snowflake Documentation: [SHOW GRANTS](#)

Question: 458

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

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

Answer: C

Explanation:

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

References:

Snowflake Documentation: [Search Optimization Service](#)

Question: 459

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 nonproduction 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 nonproduction account.

Answer: B

Explanation:

To easily transfer data from a production account to a non-production account in Snowflake within the same cloud provider region, creating a data share is the most efficient approach. Data sharing allows for live, read-only access to selected data objects from the production account to the nonproduction 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

Question: 460

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

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

Answer: C

Explanation:

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

Option C: GRANT ADD SEARCH OPTIMIZATION ON SCHEMA <schema_name> TO ROLE <role>. This command grants the specified role the ability to implement search optimization at the schema level, which is essential for enhancing search capabilities within that schema.

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

Question: 461

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

- A. Disaster recovery

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

Answer: A, D

Explanation:

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

References:

Snowflake Documentation: Data Replication

Question: 462

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

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

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

Answer: C

Explanation:

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

References:

Snowflake Documentation: Understanding Snowflake Roles

Question: 463

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

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

Answer: A

Explanation:

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

References:

Snowflake Documentation: Understanding Fail-safe

Question: 464

When floating-point number columns are unloaded to CSV or JSON files, Snowflake truncates the values to approximately what?

- A. (12,2)
- B. (10,4)
- C. (14,8)
- D. (15,9)

Answer: D

Explanation:

When unloading floating-point number columns to CSV or JSON files, Snowflake truncates the values to approximately 15 significant digits with 9 digits following the decimal point, which can be represented as (15,9). This ensures a balance between accuracy and efficiency in representing floating-point numbers in text-based formats, which is essential for data interchange and processing applications that consume these files.

References:

Snowflake Documentation: Data Unloading Considerations

Question: 465

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

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

Answer: B

Explanation:

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

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

Question: 466

Which URL provides access to files in Snowflake without authorization?

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

Answer: C

Explanation:

A Pre-signed URL provides access to files stored in Snowflake without requiring authorization at the time of access. This feature allows users to generate a URL with a limited validity period that grants temporary access to a file in a secure manner. It's particularly useful for sharing data with external parties or applications without the need for them to authenticate directly with Snowflake.

References:

Snowflake Documentation: Using Pre-signed URLs

Question: 467

What does the TableScan operator represent in the Query Profile?

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

Answer: A

Explanation:

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

References:

Snowflake Documentation: Understanding the Query Profile

Question: 468

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

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

Answer: D

Explanation:

The VARIANT data type in Snowflake is used to store JSON key-value pairs along with other semistructured 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

Question: 469

How does Snowflake reorganize data when it is loaded? (Select TWO).

- A. Binary format
- B. Columnar format
- C. Compressed format
- D. Raw format
- E. Zipped format

Answer: B, C

Explanation:

When data is loaded into Snowflake, it undergoes a reorganization process where the data is stored in a columnar format and compressed. The columnar storage format enables efficient querying and data retrieval, as it allows for reading only the necessary columns for a query, thereby reducing IO operations. Additionally, Snowflake uses advanced compression techniques to minimize storage costs and improve performance. This combination of columnar storage and compression is key to Snowflake's data warehousing capabilities.

References:

Snowflake Documentation: Data Storage and Organization

Question: 470

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

- A. Create an account for an organization.
- B. Edit the account data for an organization.
- C. Delete the account data for an organization.

- D. View usage information for all accounts in an organization.
- E. Select all the data in tables for all accounts in an organization.

Answer: A, D

Explanation:

The ORGADMIN role in Snowflake is an organizational-level role that provides administrative capabilities across the entire organization, rather than being limited to a single Snowflake account. Users with this role can:

A . Create an account for an organization: The ORGADMIN role has the privilege to create new Snowflake accounts within the organization, allowing for the expansion and management of the organization's resources.

D . View usage information for all accounts in an organization: This role also has access to comprehensive usage and activity data across all accounts within the organization. This is crucial for monitoring, cost management, and optimization at the organizational level.

References:

Snowflake Documentation: Understanding Role-Based Access Control

Question: 471

How are network policies defined in Snowflake?

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

Answer: D

Explanation:

Network policies in Snowflake are defined as a set of rules that manage the network-level access to

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

References:

Question: 472

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

Explanation:

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.

Question: 473

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

Explanation:

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

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

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

Options B, D, and E do not directly relate to observing the effects of query pruning. "Partitions total" shows the total available, not the impact of pruning, while "Bytes read from result" and "Bytes written" relate to output rather than the efficiency of data scanning.

References: Snowflake documentation on performance tuning and query optimization techniques, specifically how query pruning affects data access.

Question: 474

What is the default value in the Snowflake Web Interface (UI) for auto suspending a Virtual Warehouse?

- A. 1 minutes
- B. 5 minutes
- C. 10 minutes
- D. 15 minutes

Answer: C

Explanation:

The default value for auto-suspending a Virtual Warehouse in the Snowflake Web Interface (UI) is 10 minutes. This setting helps manage compute costs by automatically suspending warehouses that are not in use, ensuring that compute resources are efficiently allocated and not wasted on idle warehouses.

References:

Snowflake Documentation: Virtual Warehouses

Question: 475

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

Explanation:

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

Question: 476

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

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

Answer: D

Explanation:

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

References:

Snowflake Documentation: Semi-structured Data Types

Question: 477

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

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

E. Infrastructure management

Answer: D, E

Explanation:

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

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

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

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

References:

Snowflake Documentation: Overview of Snowflake Cloud Services

Question: 478

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

Explanation:

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

Question: 479

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

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

Answer: C

Explanation:

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

References:

Snowflake Documentation: Query Performance Optimization

Question: 480

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

- A. Use the SHOW STAGES command.
- B. Use the DESCRIBE STAGE command.
- C. Use the GET_ABSOLUTE_PATH function.
- D. Use the GET_STAGE_LOCATION function.

Answer: C

Explanation:

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

References:

Snowflake Documentation: Working with Stages

Question: 481

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

A. UTF-8

B. UTF-16

C. ISO-8859-1

D. WINDOWS-1253

Answer: A

Explanation:

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

References:

Snowflake Documentation: Data Loading and Unloading

Question: 482

Which statement accurately describes Snowflake's architecture?

A. It uses a local data repository for all compute nodes in the platform.

B. It is a blend of shared-disk and shared-everything database architectures.

C. It is a hybrid of traditional shared-disk and shared-nothing database architectures.

D. It reorganizes loaded data into internal optimized, compressed, and row-based format.

Answer: C

Explanation:

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

References:

Snowflake Documentation: Snowflake Architecture

Question: 483

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

Explanation:

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's 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

Question: 484

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

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

Answer: B C

Explanation:

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

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

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

References:

Snowflake Documentation: Transient Tables

Question: 485

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

Explanation:

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:

Snowflake Documentation: Virtual Warehouses

Question: 486

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

Explanation:

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

Question: 487

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

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

Answer: A

Explanation:

The Fail-safe period for a transient table in Snowflake, regardless of the edition (including Enterprise edition and higher), is 0 days. Fail-safe is a data protection feature that provides additional retention beyond the Time Travel period for recovering data in case of accidental deletion or corruption.

However, transient tables are designed for temporary or short-term use and do not benefit from the Fail-safe feature, meaning that once their Time Travel period expires, data cannot be recovered.

References:

Snowflake Documentation: Understanding Fail-safe

Question: 488

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

- A. TABLES
- B. TABLE_STORAGE_METRICS
- C. STORAGE_DAILY_HISTORY
- D. STORAGE_USAGE

Answer: B

Explanation:

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

References:

Snowflake Documentation: TABLE_STORAGE_METRICS View

Question: 489

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

Explanation:

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

Question: 490

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

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

Answer: D

Explanation:

The worksheet and database explorer feature in Snowsight allows users to tag frequently accessed worksheets for ease

of access. This functionality helps users organize and quickly navigate to the worksheets they use most often, enhancing productivity and streamlining the data exploration and analysis process within Snowsight, Snowflake's web-based query and visualization interface.

References:

Snowflake Documentation: Snowsight (UI for Snowflake)

Question: 491

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

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

Answer: A, E

Explanation:

Characteristics of reader accounts in Snowflake include:

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

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

References:

Snowflake Documentation: Reader Accounts

Question: 492

A Snowflake user wants to temporarily bypass a network policy by configuring the user object property `MINS_TO_BYPASS_NETWORK_POLICY`.

What should they do?

- A. Use the SECURITYADMIN role.

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

Answer: C

Explanation:

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

References:

Snowflake Documentation: User Management

Question: 493

When using the `ALLOW_CLIENT_MFA_CACHING` parameter, how long is a cached Multi-Factor Authentication (MFA) token valid for?

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

Answer: C

Explanation:

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

Question: 494

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

Explanation:

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

Question: 495

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

A. Grant the additional privileges to a custom role.

B. Grant the additional privileges to the ACCOUNTADMIN role.

C. Grant the additional privileges to the SYSADMIN role.

D. Grant the additional privileges to the ORGADMIN role.

Answer: A

Explanation:

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

References:

Snowflake Documentation: Roles and Privileges

Question: 496

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

A. ARRAY_AGG

- B. OBJECT_AGG
- C. ARRAY_CONSTRUCT
- D. OBJECT_CONSTRUCT

Answer: D

Explanation:

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

References:

Snowflake Documentation: Semi-structured Data Functions

Question: 497

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: B D

Explanation:

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

Question: 498

Which Snowflake layer is associated with virtual warehouses?

- A. Cloud services
- B. Query processing
- C. Elastic memory
- D. Database storage

Answer: B

Explanation:

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 costeffectiveness.

References:

Snowflake Documentation: Snowflake Architecture

Question: 499

What is it called when a customer managed key is combined with a Snowflake managed key to create a composite key for encryption?

- A. Hierarchical key model
- B. Client-side encryption
- C. Tri-secret secure encryption
- D. Key pair authentication

Answer: C

Explanation:

Tri-secret secure encryption is a security model employed by Snowflake that involves combining a customer-managed

key with a Snowflake-managed key to create a composite key for encrypting data. This model enhances data security by requiring both the customer-managed key and the Snowflake-managed key to decrypt data, thus ensuring that neither party can access the data independently. It represents a balanced approach to key management, leveraging both customer control and Snowflake's managed services for robust data encryption.

References:

Snowflake Documentation: Encryption and Key Management

Question: 500

What should be used when creating a CSV file format where the columns are wrapped by single quotes or double quotes?

- A. BINARY_FORMAT
- B. ESCAPE_UNENCLOSED_FIELD
- C. FIELD_OPTIONALLY_ENCLOSED_BY
- D. SKIP_BYTE_ORDER_MARK

Answer: C

Explanation:

When creating a CSV file format in Snowflake and the requirement is to wrap columns by single quotes or double quotes, the FIELD_OPTIONALLY_ENCLOSED_BY parameter should be used in the file format specification. This parameter allows you to define a character (either a single quote or a double quote) that can optionally enclose each field in the CSV file, providing flexibility in handling fields that contain special characters or delimiters as part of their data.

References:

Snowflake Documentation: CSV File Format

Question: 501

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

Explanation:

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

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

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

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

References:

Snowflake Documentation: Window Functions

Question: 502

How does Snowflake describe its unique architecture?

- A. A single-cluster shared data architecture using a central data repository and massively parallel processing (MPP)
- B. A multi-duster shared nothing architecture using a soloed data repository and massively parallel processing (MPP)
- C. A single-cluster shared nothing architecture using a sliced data repository and symmetric multiprocessing (SMP)
- D. A multi-cluster shared nothing architecture using a siloed data repository and symmetric multiprocessing (SMP)

Answer: A

Explanation:

Snowflake's unique architecture is described as a multi-cluster, shared data architecture that leverages massively parallel processing (MPP). This architecture separates compute and storage resources, enabling Snowflake to scale them independently. It does not use a single cluster or rely solely on symmetric multiprocessing (SMP); rather, it uses a combination of shared-nothing architecture for compute clusters (virtual warehouses) and a centralized storage layer for data, optimizing for both performance and scalability.

References:

Snowflake Documentation: Snowflake Architecture Overview

Question: 503

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

Explanation:

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

Question: 504

By default, how long is the standard retention period for Time Travel across all Snowflake accounts?

- A. 0 days
- B. 1 day

C. 7 days

D. 14 days

Answer: B

Explanation:

By default, the standard retention period for Time Travel in Snowflake is 1 day across all Snowflake accounts. Time Travel enables users to access historical data within this retention window, allowing for point-in-time data analysis and recovery. This feature is a significant aspect of Snowflake's data management capabilities, offering flexibility in handling data changes and accidental deletions.

References:

Snowflake Documentation: Using Time Travel

Question: 505

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

A. ACCOUNTADMIN

B. SECURITYADMIN

C. SYSADMIN

D. ORGADMIN

Answer: A

Explanation:

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

References:

Snowflake Documentation: Roles

Question: 506

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

A. The base tables do not change frequently.

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

Answer: A

Explanation:

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

References:

Snowflake Documentation: Materialized Views

Question: 507

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

- A. LIST
- B. ALTER FILE FORMAT
- C. DESCRIBE FILE FORMAT
- D. SHOW FILE FORMATS

Answer: D

Explanation:

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

References:

Snowflake Documentation: SHOW FILE FORMATS

Question: 508

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

- A. Syntax error

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

Answer: B, D

Explanation:

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

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

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

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

References:

Snowflake Documentation: Query Profile

Top of Form

Question: 509

By default, which role has access to the SYSTEM\$GLOBAL_ACCOUNT_SET_PARAMETER function?

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

Answer: A

Explanation:

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](#)

Question: 510

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

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

Answer: C

Explanation:

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

References:

Snowflake Documentation: [Virtual Warehouses Billing](#)

Question: 511

Which Snowflake object does not consume and storage costs?

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

Answer: C

Explanation:

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

Question: 512

What type of function returns one value for each invocation?

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

Answer: B

Explanation:

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

Question: 513

For Directory tables, what stage allows for automatic refreshing of metadata?

- A. User stage
- B. Table stage
- C. Named internal stage
- D. Named external stage

Answer: D

Explanation:

For directory tables, a named external stage allows for the automatic refreshing of metadata. This capability is particularly useful when dealing with files stored on external storage services (like Amazon S3, Google Cloud Storage, or Azure Blob Storage) and accessed through Snowflake. The external stage references these files, and the directory table's metadata can be automatically updated to reflect changes in the underlying files.

References:

Snowflake Documentation: External Stages

Question: 514

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

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

Answer: B

Explanation:

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

References:

Snowflake Documentation: REVOKE ROLE

Question: 515

The VALIDATE table function has which parameter as an input argument for a Snowflake user?

- A. Last_QUERY_ID

B. CURRENT_STATEMENT

C. UUID_STRING

D. JOB_ID

Answer: C

Explanation:

The VALIDATE table function in Snowflake would typically use a unique identifier, such as a UUID_STRING, as an input argument. This function is designed to validate the data within a table against a set of constraints or conditions, often requiring a specific identifier to reference the particular data or job being validated.

References:

There is no direct reference to a VALIDATE table function with these specific parameters in Snowflake documentation. It seems like a theoretical example for understanding function arguments. Snowflake documentation on UDFs and system functions can provide guidance on how to create and use custom functions for similar purposes.

Question: 516

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

A. FLATTEN

B. CHECK_JSON

C. PARSE_JSON

D. TO_VARIANT

Answer: C

Explanation:

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

References:

Snowflake Documentation: PARSE_JSON

Question: 517

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

Explanation:

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

Question: 518

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

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

Answer: C

Explanation:

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

References:

Snowflake Documentation: Materialized Views

Question: 519

What are characteristic of Snowsight worksheet? (Select TWO.)

- A. Worksheets can be grouped under folder, and a folder of folders.
- B. Each worksheet is a unique Snowflake session.
- C. Users are limited to running only one on a worksheet.
- D. The Snowflake session ends when a user switches worksheets.
- E. Users can import worksheets and share them with other users.

Answer: A, E

Explanation:

Characteristics of Snowsight worksheets in Snowflake include:

A . Worksheets can be grouped under folders, and a folder of folders: This organizational feature allows users to efficiently manage and categorize their worksheets within Snowsight, Snowflake's web-based UI, enhancing the user experience by keeping related worksheets together.

E . Users can import worksheets and share them with other users: Snowsight supports the sharing of worksheets among users, fostering collaboration by allowing users to share queries, analyses, and

findings. This feature is crucial for collaborative data exploration and analysis workflows.

References:

Snowflake Documentation: Snowsight (UI for Snowflake)

Question: 520

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

How will those object names be resolved during execution?

- A. Snowflake will resolve them according to the SEARCH_PATH parameter.

- B. Snowflake will only check the schema the UDF belongs to.
- C. Snowflake will first check the current schema, and then the schema the previous query used
- D. Snowflake will first check the current schema, and then the PUBLIC schema of the current database.

Answer: D

Explanation:

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

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>

Question: 521

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

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

Answer: D

Explanation:

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

References:

Snowflake Documentation: Using Materialized Views

Question: 522

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

Explanation:

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

Question: 523

What happens when a network policy includes values that appear in both the allowed and blocked IP address list?

- A. Those IP addresses are allowed access to the Snowflake account as Snowflake applies the allowed IP address list first.
- B. Those IP addresses are denied access to the Snowflake account as Snowflake applies the blocked IP address list first.
- C. Snowflake issues an alert message and adds the duplicate IP address values to both the allowed and blocked IP address lists.
- D. Snowflake issues an error message and adds the duplicate IP address values to both the allowed and blocked IP address list.

Answer: B

Explanation:

In Snowflake, when setting up a network policy that specifies both allowed and blocked IP address lists, if an IP

address appears in both lists, access from that IP address will be denied. The reason is that Snowflake prioritizes security, and the presence of an IP address in the blocked list indicates it should not be allowed regardless of its presence in the allowed list. This ensures that access controls remain stringent and that any potentially unsafe IP addresses are not inadvertently permitted access.

References:

Snowflake Documentation: Network Policies

Question: 524

What is the MINIMUM permission needed to access a file URL from an external stage?

- A. MODIFY
- B. READ
- C. SELECT
- D. USAGE

Answer: D

Explanation:

To access a file URL from an external stage in Snowflake, the minimum permission required is USAGE on the stage object. USAGE permission allows a user to reference the stage in SQL commands, necessary for actions like listing files or loading data from the stage, but does not permit the user to alter or drop the stage.

References:

Snowflake Documentation: Access Control

Question: 525

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

- A. SYSTEMADMIN_TAG
- B. SYSTEM\$GET_PRIVATELINK
- C. SYSTEMSALLONTLIST
- D. SYSTEMAUTHORIZE

Answer: B

Explanation:

The SYSTEM\$GET_PRIVATELINK function in Snowflake provides proxy information necessary for configuring PrivateLink

connections, which can protect Snowsight as well as other Snowflake services. PrivateLink enhances security by allowing Snowflake to be accessed via a private connection within a cloud provider's network, reducing exposure to the public internet.

References:

Snowflake Documentation: PrivateLink Setup

Question: 526

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

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

Answer: A

Explanation:

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

References:

Snowflake Documentation: COPY INTO Location

Topic 6, Exam pool F

Question: 527

What type of account can be used to share data with a consumer who does not have a Snowflake account?

- A. Data provider
- B. Data consumer
- C. Reader
- D. Organization

Answer: C

Explanation:

A Reader account in Snowflake can be used to share data with a consumer who does not have a Snowflake account. Reader accounts are a type of shared account provided by data providers to external data consumers, allowing them to access and query shared data using Snowflake's web interface without needing their own Snowflake account.

References:

Snowflake Documentation: Reader Accounts

Question: 528

Which privilege is needed for a Snowflake user to see the definition of a secure view?

- A. OWNERSHIP
- B. MODIFY
- C. CREATE
- D. USAGE

Answer: A

Explanation:

To see the definition of a secure view in Snowflake, the minimum privilege required is OWNERSHIP of the view. Ownership grants the ability to view the definition as well as to modify or drop the view. Secure views are designed to protect sensitive data, and thus the definition of these views is restricted to users with sufficient privileges to ensure data security.

References:

Snowflake Documentation: Secure Views

Question: 529

Which function should be used to find the query ID of the second query executed in a current session?

- A. Select LAST_QUERY_ID(-2)
- B. Select LAST_QUERY_ID(2)
- C. Select LAST_QUERY_ID(1)
- D. Select LAST_QUERY_ID(2)

Answer: A

Explanation:

The correct function to find the query ID of the second query executed in the current session is `SELECT LAST_QUERY_ID(-2)`. The `LAST_QUERY_ID` function returns the query ID for the most recent query executed in the session when called with no arguments. When used with an argument, it can retrieve the ID of previous queries within the same session, where `-2` would reference the second most recent query executed.

References:

There's a clarification needed here; Snowflake's documentation indicates `LAST_QUERY_ID()` function does not accept arguments. It returns the ID of the last query executed in the session. To find the query ID of the second last executed query, users typically need to track query IDs manually or use session history views.

Question: 530

Which command will unload data from a table into an external stage?

- A. PUT
- B. INSERT
- C. COPY INTO <location>
- D. GET

Answer: C

Explanation:

In Snowflake, the `COPY INTO <location>` command is used to unload (export) data from a Snowflake table to an external stage, such as an S3 bucket, Azure Blob Storage, or Google Cloud Storage. This command allows users to specify the format, file size, and other options for the data being unloaded, making it a flexible solution for exporting data from Snowflake to external storage solutions for further use or analysis.

References: Snowflake Documentation on Data Unloading

Question: 531

Which system-defined Snowflake role has permission to rename an account and specify whether the original URL can be used to access the renamed account?

- A. ACCOUNTADMIN

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

Answer: A

Explanation:

The ACCOUNTADMIN role in Snowflake has the highest level of privileges, including the ability to manage accounts, users, roles, and all objects within the account. This role is specifically granted the permission to rename an account and specify whether the original URL can be used to access the renamed account. The ACCOUNTADMIN role encompasses broad administrative capabilities, ensuring that users assigned this role can perform critical account management tasks.

References: Snowflake Documentation on Roles and Permissions

Question: 532

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

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

Answer: A, D

Explanation:

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

References: Snowflake Documentation on Table Types

Question: 533

What is a non-configurable feature that provides historical data that Snowflake may recover during a 7-day period?

- A. Fail-safe
- B. Time Travel
- C. Cloning
- D. Account replication

Answer: A

Explanation:

Fail-safe is a non-configurable feature in Snowflake that provides an additional layer of data protection beyond Time Travel. Time Travel allows users to access historical data within a configurable period (up to 90 days), while Fail-safe provides an additional 7-day period during which Snowflake retains historical data to recover from significant data loss or corruption incidents. This period is not accessible by users but can be used by Snowflake support to assist in data recovery efforts.

References: Snowflake Documentation on Fail-safe and Time Travel

Question: 534

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: A, E

Explanation:

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

Question: 535

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: CE

Explanation:

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>

Question: 536

What does Snowflake recommend for a user assigned the ACCOUNTADMIN role?

- A. The ACCCUKTMKIN role should be set as tie user's default role.
- B. The user should use federated authentication instead of a password
- C. The user should be required to use Multi-Factor Authentication (MFA).
- D. There should be just one user with the ACCOUNTADMIN role in each Snowflake account.

Answer: C

Explanation:

For users assigned the ACCOUNTADMIN role, Snowflake recommends enforcing Multi-Factor Authentication (MFA) to enhance security. The ACCOUNTADMIN role has extensive permissions, making it crucial to secure accounts held by such users against unauthorized access. MFA adds an additional layer of security by requiring a second form of verification beyond just the username and password, significantly reducing the risk of account compromise.

References: Snowflake Security Best Practices

Question: 537

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

Explanation:

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

Question: 538

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

Explanation:

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

References: Snowflake Documentation on Time Travel

Question: 539

Which table function should be used to view details on a Directed Acyclic Graph (DAG) run that is presently scheduled or is executing?

- A. TASK_HISTORY
- B. TASK_DEPENDENTS
- C. CURRENT_TASK_GRAPHS
- D. COMPLETE_TASK_GRAPHS

Answer: C

Explanation:

The CURRENT_TASK_GRAPHS table function is designed to provide information on Directed Acyclic Graphs (DAGs) that are currently scheduled or executing within Snowflake. This function offers insights into the structure and status of task chains, enabling users to monitor and troubleshoot task executions. DAGs in Snowflake represent sequences of tasks with dependencies, and understanding their current state is crucial for managing complex workflows.

References: Snowflake Documentation on Task Management

Question: 540

A Snowflake user is trying to load a 125 GB file using SnowSQL. The file continues to load for almost an entire day. What will happen at the 24 hour mark?

- A. The file will continue to load until all contents are loaded.
- B. The file will stop loading and all data up to that point will be committed.
- C. The file loading could be aborted without any portion of the file being committed.
- D. The file's number of allowable hours to load can be programmatically controlled to load easily into Snowflake

Answer: C

Explanation:

When attempting to load large files, such as a 125 GB file, into Snowflake using SnowSQL, the process might encounter limitations related to the maximum execution time for queries or data loading operations. If the loading process exceeds this time limit (typically around 24 hours), it could be aborted without committing any part of the file to the database.

This behavior is designed to prevent indefinite resource consumption and to maintain system stability, emphasizing the need for optimizing data load operations, possibly through file segmentation or parallel loading strategies. References:

Snowflake Documentation on Data Loading Considerations

Question: 541

Which user preferences can be set for a user profile in Snowsight? (Select TWO).

- A. Multi-Factor Authentication (MFA)
- B. Default database
- C. Default schema
- D. Notification
- E. Username

Answer: B, C

Explanation:

In Snowsight, Snowflake's web interface, user preferences can be customized to enhance the user experience. Among these preferences, users can set a default database and default schema. These settings streamline the user experience by automatically selecting the specified database and schema when the user initiates a new session or query, reducing the need to manually specify these parameters for each operation. This feature is particularly useful for users who frequently work within a specific database or schema context.

References: Snowflake Documentation on Snowsight User Preferences

Question: 542

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: D

Explanation:

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.

Question: 543

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

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

Answer: D

Explanation:

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

Question: 544

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

Explanation:

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

Question: 545

How does Snowflake define its approach to Discretionary Access Control (DAC)?

- A. A defined level of access to an object
- B. An entity in which access can be granted
- C. Each object has an owner, who can in turn grant access to that object.
- D. Access privileges are assigned to roles, which are in turn assigned to users

Answer: D

Explanation:

Snowflake implements Discretionary Access Control (DAC) by using a role-based access control model. In this model, access privileges are not directly assigned to individual objects or users but are encapsulated within roles. These roles are then assigned to users, effectively granting them the access privileges contained within the role. This approach allows for granular control over database access, making it easier to manage permissions in a scalable and flexible manner. References: Snowflake Documentation on Access Control

Question: 546

How many credits does a size 3X-Large virtual warehouse consume if it runs continuously for 2 hours?

- A. 32
- B. 64
- C. 128
- D. 256

Answer: C

Explanation:

In Snowflake, the consumption of credits by a virtual warehouse is determined by its size and the duration for which it runs. A size 3X-Large virtual warehouse consumes 128 credits if it runs continuously for 2 hours. This consumption rate is based on the principle that larger warehouses, capable of providing greater computational resources and throughput, consume more credits per hour of operation. The specific rate of consumption is defined by Snowflake's

pricing model and the scale of the virtual warehouse.

References: Snowflake Pricing Documentation

Question: 547

Which command is used to take away staged files from a Snowflake stage after a successful data ingestion?

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

Answer: C

Explanation:

The REMOVE command is used in Snowflake to delete files from a stage after they have been successfully ingested into Snowflake tables. This command helps manage storage by allowing users to clean up staged files that are no longer needed, ensuring that the stage does not accumulate unnecessary data over time. Unlike DELETE, DROP, or TRUNCATE commands, which are used for managing data within Snowflake tables or dropping objects, REMOVE specifically targets the management of files in stages.

References: Snowflake Documentation on Stages and File Management

Question: 548

Which Snowflake table supports unstructured data?

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

Answer: D

Explanation:

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

Question: 549

If a virtual warehouse is suspended, what happens to the warehouse cache?

- A. The cache is dropped when the warehouse is suspended and is no longer available upon restart.
- B. The warehouse cache persists for as long the warehouse exists, regardless of its suspension status.
- C. The cache is maintained for up to two hours and can be restored if the warehouse is restarted within this limit.
- D. The cache is maintained for the auto suspend duration and can be restored if the warehouse is restarted within this limit.

Answer: A

Explanation:

When a virtual warehouse in Snowflake is suspended, the cache is dropped and is no longer available upon restart. This means that all cached data, including results and temporary data, are cleared from memory. The purpose of this behavior is to conserve resources while the warehouse is not active. Upon restarting the warehouse, it will need to reload any data required for queries from storage, which may result in a slower initial performance until the cache is repopulated. This is a critical consideration for managing performance and cost in Snowflake.

Question: 550

Who can create network policies within Snowflake? (Select TWO).

- A. SYSADMIN only
- B. ORCADMIN only
- C. SECURITYADMIN or higher roles
- D. A role with the CREATE NETWORK POLICY privilege
- E. A role with the CREATE SECURITY INTEGRATION privilege

Answer: C, D

Explanation:

In Snowflake, network policies define the allowed IP address ranges from which users can connect to Snowflake, enhancing security by restricting access based on network location. The creation and management of network policies

require sufficient privileges. Specifically, a user with the SECURITYADMIN role or any role with higher privileges, such as ACCOUNTADMIN, can create network policies. Additionally, a custom role can be granted the CREATE NETWORK POLICY privilege, enabling users assigned to that role to also create network policies. This approach allows for flexible and secure management of network access to Snowflake.

References: Snowflake Documentation on Network Policies

Question: 551

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

Explanation:

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

Question: 552

A stream can be created on which Snowflake objects to record data manipulation language(DML) changes? (Select TWO).

- A. Database
- B. Standard tables
- C. Standard tables
- D. Standard views
- E. Schemas
- F. Pipes

Answer: B, C

Explanation:

Snowflake streams are objects that enable users to track changes (inserts, updates, and deletes) to the data in tables, facilitating real-time data processing and integration workflows. Streams can be created on standard tables, capturing DML changes made to these tables so that downstream processes can consume the changes incrementally. This feature supports efficient data ETL, replication, and real-time analytics by providing a mechanism to process only the data that has changed. Note: The correct options should likely include a distinction between "Standard tables" and another object type such as "External tables" rather than repeating "Standard tables" twice. References: Snowflake Documentation on Streams

Question: 553

What can a user with a reader account do in Snowflake?

- A. Load new data
- B. Update existing data
- C. Create a new share
- D. Query shared data

Answer: D

Explanation:

In Snowflake, a user with a reader account primarily has read-only access to the shared data. Reader accounts are created to enable third parties or separate business units to access and query data shared with them without allowing them to modify the underlying data. This means a user with a reader account can perform operations like querying shared data to analyze or report on it but cannot load new data, update existing data, or create new shares. This setup is crucial for maintaining data governance and security while enabling data sharing and collaboration. References: Snowflake Documentation on Reader Accounts

Question: 554

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

Explanation:

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

Question: 555

Snowflake users can create a resource monitor at which levels? (Select TWO).

- A. User level
- B. Pipe level
- C. Account level
- D. Cloud services level
- E. Virtual warehouse level

Answer: C, E

Explanation:

Resource monitors in Snowflake are tools used to track and control the consumption of compute resources, ensuring that usage stays within defined limits. These monitors can be created at the account level, allowing administrators to set overall resource consumption limits for the entire Snowflake account. Additionally, resource monitors can be set at the virtual warehouse level, enabling more granular control over the resources consumed by individual warehouses. This dual-level capability allows organizations to manage their Snowflake usage efficiently, preventing unexpected costs and optimizing performance.

References: Snowflake Documentation on Resource Monitors

Top of Form

Question: 556

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

Explanation:

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

Question: 557

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

Explanation:

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

Question: 558

Which roles can make grant decisions to objects within a managed access schema? (Select TWO)

- A. ACCOUNTADMIN
- B. SECURITYADMIN
- C. SYSTEMADMIN
- D. ORGADMIN
- E. USERADMIN

Answer: A, B

Explanation:

Managed Access Schemas: These are a special type of schema designed for fine-grained access control in Snowflake.

Roles with Grant Authority:

ACCOUNTADMIN: The top-level administrative role can grant object privileges on all objects within the account, including managed access schemas.

SECURITYADMIN: Can grant and revoke privileges on objects within the account, including managed access schemas.

Important Note: The ORGADMIN role focuses on organization-level management, not object access control.

Question: 559

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

Explanation:

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

Question: 560

What are the main differences between the account usage views and the information schema views? (Select TWO).

- A. No active warehouse is needed to query account usage views but one is needed to query information schema views.
- B. Account usage views do not contain data about tables but information schema views do.
- C. Account usage views contain dropped objects but information schema views do not.
- D. Data retention for account usage views is 1 year but is 7 days to 6 months for information schema views, depending on the view.
- E. Information schema views are read-only but account usage views are not.

Answer: C, D

Explanation:

The account usage views in Snowflake provide historical usage data about the Snowflake account, and they retain this data for a period of up to 1 year. These views include information about dropped objects, enabling audit and tracking activities. On the other hand, information schema views provide metadata about database objects currently in use, such as tables and views, but do not include dropped objects. The retention of data in information schema views varies, but it is generally shorter than the retention for account usage views, ranging from 7 days to a maximum of 6 months, depending on the specific view.

References: Snowflake Documentation on Account Usage and Information Schema

Question: 561

In which hierarchy is tag inheritance possible?

- A. Organization » Account » Role
- B. Account » User » Schema
- C. Database » View » Column
- D. Schema » Table » Column

Answer: D

Explanation:

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

Question: 562

Which data types can be used in a Snowflake table that holds semi-structured data? (Select TWO).

- A. ARRAY
- B. BINARY
- C. TEXT
- D. VARIANT
- E. VARCHAR

Answer: A, D

Explanation:

Snowflake supports semi-structured data types that can store complex data structures within a single column. The VARIANT data type can hold structured and semi-structured data formats such as JSON, Avro, ORC, Parquet, and XML. The ARRAY data type is also used for semi-structured data and can store an ordered list of elements. These data types enable users to work with semi-structured data directly in Snowflake without needing to flatten the data into a relational schema first.

References: Snowflake Documentation on Semi-Structured Data Types

Question: 563

Which Snowflake database object can be shared with other accounts?

- A. Tasks
- B. Pipes
- C. Secure User-Defined Functions (UDFs)
- D. Stored Procedures

Answer: C

Explanation:

In Snowflake, Secure User-Defined Functions (UDFs) can be shared with other accounts using Snowflake's data sharing feature. This allows different Snowflake accounts to securely execute the UDFs without having direct access to the underlying data the functions operate on, ensuring privacy and security. The sharing is facilitated through shares created

in Snowflake, which can contain Secure UDFs along with other database objects like tables and views.

References: Snowflake Documentation on Data Sharing and Secure UDFs

Question: 564

What does Snowflake attempt to do if any of the compute resources for a virtual warehouse fail to provision during start-up?

- A. Repair the failed resources.
- B. Restart failed resources.
- C. Queue the failed resources
- D. Provision the failed resources

Answer: B

Explanation:

If any compute resources for a virtual warehouse fail to provision during startup, Snowflake will attempt to restart those failed resources. The system is designed to automatically handle transient issues that might occur during the provisioning of compute resources. By restarting the failed resources, Snowflake aims to ensure that the virtual warehouse has the necessary compute capacity to handle the user's workloads without manual intervention.

References: Snowflake Documentation on Virtual Warehouses

Question: 565

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

Explanation:

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

Question: 566

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

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

Answer: B

Explanation:

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

Question: 567

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

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

Explanation:

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,

Question: 568

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

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

Answer: D

Explanation:

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

References: Snowflake Documentation on Client Drivers

Question: 569

While running a query on a virtual warehouse in auto-scale mode, additional clusters are started immediately if which setting is configured?

- A. MAX_CLUSTER_COUNT is increased and new_max_clusters is greater than running_clusters
- B. MAX_CLUSTER_COUNT is decreased and new_max_clusters is less than running_clusters
- C. MIN_CLUSTER_COUNT is increased and new_min_clusters is greater than running_clusters
- D. MIN_CLUSTER_COUNT is decreased and new_min_clusters is less than running_clusters

Answer: C

Explanation:

Question: 570

How should a Snowflake user configure a virtual warehouse to be in Maximized mode?

- A. Set the WAREHOUSES_SIZE to 6XL.
- B. Set the STATEMENT_TIMEOUT_1M_SECONDS to 0.
- C. Set the MAX_CONCURRENCY_LEVEL to a value of 12 or larger.
- D. Set the same value for both MIN_CLUSTER_COUNT and MAX_CLUSTER_COUNT.

Answer: D

Explanation:

In Snowflake, configuring a virtual warehouse to be in a "Maximized" mode implies maximizing the resources allocated to the warehouse for its duration. This is done to ensure that the warehouse has a consistent amount of compute resources available, enhancing performance for workloads that require a high level of parallel processing or for handling high query volumes.

To configure a virtual warehouse in maximized mode, you should set the same value for both MIN_CLUSTER_COUNT and MAX_CLUSTER_COUNT. This configuration ensures that the warehouse operates with a fixed number of clusters, thereby providing a stable and maximized level of compute resources.

Reference to Snowflake documentation on warehouse sizing and scaling:

[Warehouse Sizing and Scaling](#)

[Understanding Warehouses](#)

Question: 571

Which system-defined, read-only view displays information on column lineage that specifies how data flows from source to target in a SQL write operation?

- A. ACCESS_HISTORY
- B. LOAD_HISTORY
- C. QUERY_HISTORY
- D. COPY_HISTORY

Answer: A

Explanation:

In Snowflake, the system-defined, read-only view that displays information on column lineage, which specifies how data flows from source to target in a SQL write operation, is ACCESS_HISTORY. This view is instrumental in auditing and analyzing data access patterns, as it provides detailed insights into how and from where the data is being accessed and

manipulated within Snowflake.

Reference to Snowflake documentation on ACCESS_HISTORY:

Using Access History to Audit Data Access

Question: 572

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

Explanation:

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

Question: 573

Which view in SNOWFLAKE.ACCOUNT_USAGE shows from which IP address a user connected to Snowflak?

- A. ACCESS_HOSTORY
- B. LOGIN_HISTORY
- C. SESSIONS

D. QUERY HISTORY

Answer: B

Explanation:

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

Question: 574

Given the statement template below, which database objects can be added to a share?(Select TWO).

```
GRANT <privilege> ON <object> <object_name> To SHARE <share_name>;
```

- A. Secure functions
- B. Stored procedures
- C. Streams
- D. Tables
- E. Tasks

Answer: C, D

Explanation:

In Snowflake, shares are used to share data across different Snowflake accounts securely. When you create a share, you can include various database objects that you want to share with consumers. According to Snowflake's documentation, the types of objects that can be shared include tables, secure views, secure materialized views, and streams. Secure functions and stored procedures are not shareable objects. Tasks also cannot be shared directly. Therefore, the correct answers are streams (C) and tables (D).

To share a stream or a table, you use the GRANT statement to grant privileges on these objects to a share. The syntax for sharing a table or stream involves specifying the type of object, the object name, and the share to which you are granting access. For example:

```
GRANT SELECT ON TABLE my_table TO SHARE my_share; GRANT SELECT ON STREAM my_stream TO
```

SHARE my_share;

These commands grant the SELECT privilege on a table named my_table and a stream named my_stream to a share named my_share. This enables the consumer of the share to access these objects according to the granted privileges.

Reference: Snowflake Documentation on Shares and Database Objects (<https://docs.snowflake.com>)

Question: 575

How can an administrator check for updates (for example, SCIM API requests) sent to Snowflake by the identity provider?

- A. ACCESS_HISTORY
- B. LOAD_HISTORY
- C. QUERY_HISTORY
- D. REST EVENT HISTORY

Answer: D

Explanation:

To monitor updates, such as SCIM API requests sent to Snowflake by the identity provider, an administrator can use the REST EVENT HISTORY feature. This feature allows administrators to query historical data about REST API calls made to Snowflake, including those related to user and role management through SCIM (System for Cross-domain Identity Management).

The REST EVENT HISTORY table function returns information about REST API calls made over a specified period. It is particularly useful for auditing and monitoring purposes, especially when integrating Snowflake with third-party identity providers that use SCIM for automated user provisioning and deprovisioning.

An example query to check for SCIM API requests might look like this:

```
SELECT * FROM TABLE(information_schema.rest_event_history(date_range_start=>dateadd('hours',-1,current_timestamp())) WHERE request_type = 'SCIM');
```

This query returns details on SCIM API requests made in the last hour, including the request type, the identity provider's details, and the outcome of each request.

Reference: Snowflake Documentation on REST EVENT HISTORY (https://docs.snowflake.com/en/sql-reference/functions/rest_event_history.html)

Question: 576

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

Explanation:

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>)

Question: 577

Which Snowflake privilege is required on a pipe object to pause or resume pipes?

- A. OPERATE
- B. READ
- C. SELECT
- D. USAGE

Answer: A

Explanation:

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;
```

Reference: For more information on pipe operations and privileges, you can refer to Snowflake's official documentation on pipes and security access control: <https://docs.snowflake.com/en/sql-reference/sql/grant-privilege.html> and <https://docs.snowflake.com/en/user-guide/data-pipelines-intro.html#managing-pipes>

Question: 578

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: D

Explanation:

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

Question: 579

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

Explanation:

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.

Question: 580

Authorization to execute CREATE <object> statements comes only from which role?

- A. Primary role
- B. Secondary role
- C. Application role
- D. Database role

Answer: A

Explanation:

In Snowflake, the authorization to execute CREATE <object> statements, such as creating tables, views, databases, etc., is determined by the role currently set as the user's primary role. The primary role of a user or session specifies the set of privileges (including creation privileges) that the user has. While users can have multiple roles, only the primary role is used to determine what objects the user can create unless explicitly specified in the session.

Reference: This is based on the principle of Role-Based Access Control (RBAC) in Snowflake, where roles are used to manage access permissions. The official Snowflake documentation on

Understanding and Using Roles would be the best resource to verify this information:

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

Question: 581

Which command is used to upload data files from a local directory or folder on a client machine to an internal stage, for a specified table?

- A. GET
- B. PUT
- C. CREATE STREAM
- D. COPY INTO <location>

Answer: B

Explanation:

To upload data files from a local directory or folder on a client machine to an internal stage in Snowflake, the PUT command is used. The PUT command takes files from the local file system and uploads them to an internal Snowflake stage (or a specified stage) for the purpose of preparing the data to be loaded into Snowflake tables.

Syntax Example:

```
PUT file://<local_file_path> @<stage_name>;
```

This command is crucial for data ingestion workflows in Snowflake, especially when preparing to load data using the COPY INTO command.

Reference: Snowflake Documentation on Using the PUT Command: <https://docs.snowflake.com/en/sql-reference/sql/put.html>

Question: 582

Which file function provides a URL with access to a file on a stage without the need for authentication and authorization?

- A. GET_RELATIVE_PATH
- B. GET_PRESIGNED_URL
- C. BUILD_STAGE_FILE_URL
- D. BUILD_SCOPED_FILE_URL

Answer: B

Explanation:

The GET_PREIGNED_URL file function in Snowflake provides a URL with access to a file on a stage without requiring authentication and authorization. This is particularly useful for sharing data files stored in Snowflake stages with external parties securely and conveniently. The presigned URL generated by this function gives temporary access to the file, which expires after a specified duration.

Example usage of GET_PREIGNED_URL:

```
SELECT GET_PREIGNED_URL('<stage_name>', '<file_path>');
```

This function generates a URL that can be used to directly access a file in the stage, making it easier to share data without compromising security.

Reference: Snowflake Documentation on File and Table Functions, specifically the section on GET_PREIGNED_URL:

https://docs.snowflake.com/en/sql-reference/functions/get_presigned_url.html

Top of Form

Question: 583

Which governance feature is supported by all Snowflake editions?

- A. Object tags
- B. Masking policies
- C. Row access policies
- D. OBJECT_DEPENDENCIES View

Answer: D

Explanation:

Snowflake's governance features vary across different editions, but the OBJECT_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 finegrained 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)

Question: 584

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

Explanation:

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>)

Question: 585

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

Explanation:

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>)

Question: 586

At what level is the MIN_DATA_RETENTION_TIME_IN_DAYS parameter set?

A. Account

B. Database

C. Schema

D. Table

Answer: A

Explanation:

The MIN_DATA_RETENTION_TIME_IN_DAYS parameter is set at the Account level in Snowflake. This parameter specifies the minimum number of days Snowflake retains the historical data for time travel, which allows users to access and query data as it existed at previous points in time.

Here's how to understand and adjust this parameter:

Purpose of MIN_DATA_RETENTION_TIME_IN_DAYS: This parameter is crucial for managing data lifecycle and compliance

requirements within Snowflake. It determines the minimum time frame for which you can perform operations like restoring deleted objects or accessing historical versions of data.

Setting the Parameter: Only account administrators can set or modify this parameter. It is done at the account level, impacting all databases and schemas within the account. The setting can be adjusted based on the organization's data retention policy.

Adjusting the Parameter:

To view the current setting, use:

```
SHOW PARAMETERS LIKE 'MIN_DATA_RETENTION_TIME_IN_DAYS';
```

To change the setting, an account administrator can execute:

```
ALTER ACCOUNT SET MIN_DATA_RETENTION_TIME_IN_DAYS = <number_of_days>;
```

Reference: For more details on data retention and time travel in Snowflake, refer to the official documentation:

<https://docs.snowflake.com/en/sql-reference/account-usage/storage-usage.html#data-retention-time-travel>

Question: 587

What is the MINIMUM size of a table for which Snowflake recommends considering adding a clustering key?

- A. 1 Kilobyte (KB)
- B. 1 Megabyte (MB)
- C. 1 Gigabyte (GB)
- D. 1 Terabyte (TB)

Answer: D

Explanation:

Snowflake recommends considering adding a clustering key to a table when its size reaches 1 Terabyte (TB) or larger.

Clustering keys help optimize the storage and query performance by organizing the data in a table based on the specified columns. This is particularly beneficial for large tables where data retrieval can become inefficient without proper clustering.

Why Clustering Keys Are Important: Clustering keys ensure that data stored in Snowflake is physically ordered in a way that aligns with the most frequent access patterns, thereby reducing the amount of scanned data during queries and improving performance.

Recommendation Basis: The recommendation for tables of size 1 TB or larger is based on the observation that smaller tables generally do not benefit as much from clustering, given Snowflake's architecture. However, as tables grow in size, the benefits of clustering become more pronounced.

Implementing Clustering Keys:

To set a clustering key for a table, you can use the CLUSTER BY clause during table creation or alter an existing table to add it:

```
CREATE TABLE my_table (... ) CLUSTER BY (column1, column2);
```

Or for an existing table:

```
ALTER TABLE my_table CLUSTER BY (column1, column2);
```

Reference: For additional information on clustering in Snowflake and recommendations on when to use clustering keys, consult the Snowflake documentation on managing and optimizing data clustering: <https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html>

Question: 588

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

Explanation:

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.

Reference: For more detailed information on the JAROWINKLER_SIMILARITY function and its usage, refer to the Snowflake documentation on string functions: https://docs.snowflake.com/en/sql-reference/functions/jarowinkler_similarity.html

Question: 589

Which types of subqueries does Snowflake support? (Select TWO).

- A. Uncorrelated scalar subqueries in WHERE clauses
- B. Uncorrelated scalar subqueries in any place that a value expression can be used
- C. EXISTS, ANY / ALL, and IN subqueries in WHERE clauses: these subqueries can be uncorrelated only
- D. EXISTS, ANY / ALL, and IN subqueries in where clauses: these subqueries can be correlated only
- E. EXISTS, ANY /ALL, and IN subqueries in WHERE clauses: these subqueries can be correlated or uncorrelated

Answer: B, E

Explanation:

Snowflake supports a variety of subquery types, including both correlated and uncorrelated subqueries. The correct answers are B and E, which highlight Snowflake's flexibility in handling subqueries within SQL queries.

Uncorrelated Scalar Subqueries: These are subqueries that can execute independently of the outer query. They return a

single value and can be used anywhere a value expression is allowed, offering great flexibility in SQL queries.

EXISTS, ANY/ALL, and IN Subqueries: These subqueries are used in WHERE clauses to filter the results of the main query based on the presence or absence of matching rows in a subquery. Snowflake supports both correlated and uncorrelated versions of these subqueries, providing powerful tools for complex data analysis scenarios.

Examples and Usage:

Uncorrelated Scalar Subquery:

```
SELECT * FROM employees WHERE salary > (SELECT AVG(salary) FROM employees);
```

Correlated EXISTS Subquery:

```
SELECT * FROM orders o WHERE EXISTS (SELECT 1 FROM customer c WHERE c.id = o.customer_id AND c.region = 'North America');
```

Reference: For comprehensive details on the types of subqueries supported by Snowflake and examples of their usage, consult the Snowflake documentation on subqueries: <https://docs.snowflake.com/en/sql-reference/constructs/subqueries.html>

Question: 590

Which Snowflake data governance feature can support auditing when a user query reads column data?

- A. Access History
- B. Data classification
- C. Column-level security
- D. Object dependencies

Answer: A

Explanation:

Access History in Snowflake is a feature designed to support auditing by tracking access to data within Snowflake, including when a user's query reads column data. It provides detailed information on queries executed, including the user who ran the query, the query text, and the objects (e.g., tables, views) accessed by the query. This feature is instrumental for auditing purposes, helping organizations to monitor and audit data access for security and compliance.

Reference: Snowflake Documentation on Access History which details its use for auditing access to data: <https://docs.snowflake.com/en/user-guide/access-history.html>

Question: 591

A clustering key was defined on a table, but It is no longer needed. How can the key be removed?

- A. ALTER TABLE <TABLE NAME> PURGE CLUSTERING KEY
- B. ALTER TABLE <TABLE NAME> DELETE CLUSTERING KEY
- C. ALTER TABLE <TABLE NAME> DROP CLUSTERING KEY
- D. ALTER TABLE <TABLE NAME> REMOVE CLUSTERING KEY

Answer: C

Explanation:

To remove a clustering key that was previously defined on a table in Snowflake, the correct SQL command is ALTER TABLE <TABLE NAME> DROP CLUSTERING KEY. This command removes the existing clustering key from the table, after which Snowflake will no longer re-cluster data based on that key during maintenance operations or after data loading operations.

Reference: Snowflake Documentation on altering tables, specifically the section on modifying clustering keys:
<https://docs.snowflake.com/en/sql-reference/sql/alter-table.html#modifying-clustering-keys>

Question: 592

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

Explanation:

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

Network policies are not applied to roles but are set at the account or user level. They are not activated using ALTER DATABASE SQL commands but are managed through ALTER ACCOUNT or ALTER NETWORK POLICY commands. The

management of network policies does not exclusively require the ORGADMIN role; instead, they can be managed by users with the necessary privileges on the account.

Reference: Snowflake Documentation on Network Policies: <https://docs.snowflake.com/en/user-guide/network-policies.html>

Question: 593

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

Explanation:

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>

Question: 594

Which command can be used to list all network policies available in an account?

- A. DESCRIBE SESSION POLICY
- B. DESCRIBE NETWORK POLICY
- C. SHOW SESSION POLICIES
- D. SHOW NETWORK POLICIES

Answer: D

Explanation:

To list all network policies available in an account, the correct command is SHOW NETWORK POLICIES. Network policies in Snowflake are used to define and enforce rules for how users can connect to Snowflake, including IP whitelisting and other connection requirements. The SHOW NETWORK POLICIES command provides a list of all network policies defined within the account, along with their details.

The DESCRIBE SESSION POLICY and DESCRIBE NETWORK POLICY commands do not exist in Snowflake SQL syntax. The SHOW SESSION POLICIES command is also incorrect, as it does not pertain to the correct naming convention used by Snowflake for network policy management.

Using SHOW NETWORK POLICIES without any additional parameters will display all network policies in the account, which is useful for administrators to review and manage the security configurations pertaining to network access.

Reference: Snowflake Documentation on Network Policies (<https://docs.snowflake.com/en/sql-reference/sql/show-network-policies.html>)

Question: 595

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

Explanation:

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.

Reference: Snowflake Documentation on Secure Views (<https://docs.snowflake.com/en/user-guide/views-secure.html>)

Question: 596

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

Explanation:

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.

Reference: Snowflake Documentation on Managing Warehouse Credit Usage

(<https://docs.snowflake.com/en/user-guide/warehouses-considerations.html#managing-warehouse-credit-usage>)

Question: 597

A Snowflake user wants to optimize performance for a query that queries only a small number of rows in a table. The rows require significant processing. The data in the table does not change frequently.

What should the user do?

- A. Add a clustering key to the table.
- B. Add the search optimization service to the table.
- C. Create a materialized view based on the query.
- D. Enable the query acceleration service for the virtual warehouse.

Answer: C

Explanation:

In a scenario where a Snowflake user queries only a small number of rows that require significant processing and the data in the table does not change frequently, the most effective way to optimize performance is by creating a materialized view based on the query. Materialized views store the result of the query and can significantly reduce the computation time for queries that are executed frequently over unchanged data.

Why Materialized Views: Materialized views precompute and store the result of the query. This is especially beneficial for queries that require heavy processing. Since the data does not change frequently, the materialized view will not need to be refreshed often, making it an ideal solution for this use case.

Implementation Steps:

To create a materialized view, use the following SQL command:

```
CREATE MATERIALIZED VIEW my_materialized_view AS SELECT ... FROM my_table WHERE ...;
```

When the query is run, Snowflake uses the precomputed results from the materialized view, thus skipping the need for recalculating the data and improving query performance.

Reference: For more information on materialized views and how they can be used to optimize query performance, refer to the Snowflake documentation on materialized views: <https://docs.snowflake.com/en/user-guide/views->

materialized.html

Question: 598

To use the overwrite option on insert, which privilege must be granted to the role?

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

Answer: B

Explanation:

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;
```

Reference: For additional details on inserting data with the overwrite option and the required privileges, consult the Snowflake documentation on data loading: <https://docs.snowflake.com/en/sql-reference/sql/insert.html>

Question: 599

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

Explanation:

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

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

Creating a Transient Table:

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.

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

Question: 600

What is the default access of a securable object until other access is granted?

A. No access

B. Read access

C. Write access

D. Full access

Answer: A

Explanation:

In Snowflake, the default access level for any securable object (such as a table, view, or schema) is "No access" until explicit access is granted. This means that when an object is created, only the owner of the object and roles with the necessary privileges can access it. Other users or roles will not have any form of access to the object until it is explicitly granted.

This design adheres to the principle of least privilege, ensuring that access to data is tightly controlled and that users and roles only have the access necessary for their functions. To grant access, the owner of the object or a role with the GRANT option can use the GRANT statement to provide specific privileges to other users or roles.

For example, to grant SELECT access on a table to a specific role, you would use a command similar to:

```
GRANT SELECT ON TABLE my_table TO ROLE my_role;
```

Reference: Snowflake Documentation on Access Control (<https://docs.snowflake.com/en/user-guide/security-access-control-overview.html>)

Question: 601

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: D

Explanation:

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>)

Question: 602

How does Snowflake handle the data retention period for a table if a stream has not been consumed?

- A. The data retention period is reduced to a minimum of 14 days.
- B. The data retention period is permanently extended for the table.
- C. The data retention period is temporarily extended to the stream's offset.
- D. The data retention period is not affected by the stream consumption.

Answer: C

Explanation:

In Snowflake, the use of streams impacts how the data retention period for a table is handled, particularly in scenarios where the stream has not been consumed. The key point to understand is that Snowflake's streams are designed to capture data manipulation language (DML) changes such as INSERTS, UPDATES, and DELETES that occur on a source table. Streams maintain a record of these changes until they are consumed by a DML operation or a COPY command that references the stream.

When a stream is created on a table and remains unconsumed, Snowflake extends the data retention period of the table to ensure that the changes captured by the stream are preserved. This extension is specifically up to the point in time represented by the stream's offset, which effectively ensures that the data necessary for consuming the stream's contents is retained. This mechanism is in place to prevent data loss and ensure the integrity of the stream's data, facilitating accurate and reliable data processing and analysis based on the captured DML changes.

This behavior emphasizes the importance of managing streams and their consumption appropriately to balance between data retention needs and storage costs. It's also crucial to understand how this temporary extension of the data retention period impacts the overall management of data within Snowflake, including aspects related to data lifecycle, storage cost implications, and the planning of data consumption strategies.

References:

Snowflake Documentation on Streams: Using Streams

Snowflake Documentation on Data Retention: Understanding Data Retention

Question: 603

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

Explanation:

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

Question: 604

Which feature of Snowflake's Continuous Data Protection (CDP) has associated costs?

- A. Fail-safe
- B. Network policies
- C. End-to-end encryption
- D. Multi-Factor Authentication (MFA)

Answer: A

Explanation:

Snowflake's Continuous Data Protection (CDP) features encompass several mechanisms designed to protect data and

ensure its availability and recoverability. Among these features, the one that has associated costs is Fail-safe.

Fail-safe is an additional layer of protection that kicks in after the Time Travel period expires. While

Time Travel allows users to access historical data within a defined retention period (which can vary from 1 to 90 days depending on the Snowflake edition), Fail-safe provides a further 7 days (for a total of 7 additional days beyond the Time Travel period) during which Snowflake retains the data. This period is primarily intended for Snowflake's internal operations to recover data in the event of extreme scenarios, such as significant operational failures, and is not directly accessible by customers for data recovery purposes.

The associated costs with Fail-safe arise because Snowflake continues to store the data beyond the customer-specified Time Travel period, thereby incurring additional storage costs. It's important to note that while users do not incur direct costs for enabling Fail-safe (as it is an automatic feature of Snowflake), the extended storage of data during this period contributes to overall storage costs.

References:

Snowflake Documentation on Continuous Data Protection: Continuous Data Protection (CDP)

Snowflake Documentation on Fail-safe: Understanding Fail-safe

Question: 605

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

Explanation:

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

Question: 606

Which of the following SQL statements will list the version of the drivers currently being used?

- A. Execute SELECT CURRENT_ODBC_CLIENT(); from the Web UI
- B. Execute SELECT CURRENT_JDBC_VERSION() ; from SnowSQL
- C. Execute SELECT CURRENT_CLIENT(); from an application
- D. Execute SELECT CURRENT_VERSION (); from the Python Connector

Answer: B

Explanation:

The correct SQL statement to list the version of the JDBC (Java Database Connectivity) drivers currently being used is to execute SELECT CURRENT_JDBC_VERSION(); from within SnowSQL or any client application that utilizes JDBC for connecting to Snowflake. Snowflake provides specific functions to query the version of the client drivers or connectors being used, such as JDBC, ODBC, and others. The CURRENT_JDBC_VERSION() function is designed specifically to return the version of the JDBC driver in use.

It's important to note that Snowflake supports various types of drivers and connectors for connecting to different client applications, including ODBC, JDBC, Python, and others. Each of these connectors has its own method or function for querying the current version in use. For JDBC, the appropriate function is CURRENT_JDBC_VERSION(), reflecting the specificity required to obtain version information relevant to the JDBC driver specifically.

References:

Snowflake Documentation on Client Functions: This information can typically be found in the Snowflake documentation under the section that covers SQL functions, particularly those functions that provide information about the client or session.

Question: 607

Which object type is granted permissions for reading a table?

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

Answer: B

Explanation:

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

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

Question: 608

When should a stored procedure be created with caller's rights?

- A. When the caller needs to be prevented from viewing the source code of the stored procedure
- B. When the caller needs to run a statement that could not execute outside of the stored procedure
- C. When the stored procedure needs to run with the privileges of the role that called the stored procedure
- D. When the stored procedure needs to operate on objects that the caller does not have privileges on

Answer: C

Explanation:

Stored procedures in Snowflake can be created with either 'owner's rights' or 'caller's rights'. A stored procedure created with caller's rights executes with the privileges of the role that calls the procedure, not the privileges of the role that owns the procedure. This is particularly useful in scenarios where the procedure needs to perform operations that depend on the caller's access permissions, ensuring that the procedure can only access objects that the caller is authorized to access.

Reference: Snowflake's official documentation on stored procedures, specifically the section on execution context (caller's rights vs. owner's rights), provides detailed guidance on when to use caller's rights:

<https://docs.snowflake.com/en/sql-reference/stored-procedures-usage.html#caller-s-rights-vs-owner-s-rights>

Question: 609

Which function can be used with the copy into <location> statement to convert rows from a relational table to a single variant column, and to unload rows into a JSON file?

- A. FLATTEN
- B. OBJECT_AS
- C. OBJECT_CONSTRUCT
- D. TO VARIANT

Answer: D

Explanation:

The correct function to use with the COPY INTO <location> statement to convert rows from a relational table into a single variant column and to unload rows into a JSON file is TO VARIANT. The TO VARIANT function is used to explicitly convert a value of any supported data type into a VARIANT data type. This is particularly useful when needing to aggregate multiple columns or complex data structures into a single JSON-formatted string, which can then be unloaded into a file.

In the context of unloading data, the COPY INTO <location> statement combined with TO VARIANT enables the conversion of structured data from Snowflake tables into a semi-structured VARIANT format, typically JSON, which can then be efficiently exported and stored. This approach is often utilized for data integration scenarios, backups, or when data needs to be shared in a format that is easily consumed by various applications or services that support JSON.

References:

[Snowflake Documentation on Data Unloading: Unloading Data](#)

[Snowflake Documentation on VARIANT Data Type: Working with JSON](#)

Question: 610

By default, which role can create resource monitors?

- A. ACCOUNTADMIN

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

Answer: A

Explanation:

The role that can by default create resource monitors in Snowflake is the ACCOUNTADMIN role. Resource monitors are a crucial feature in Snowflake that allows administrators to track and control the consumption of compute resources, ensuring that usage stays within specified limits. The creation and management of resource monitors involve defining thresholds for credits usage, setting up notifications, and specifying actions to be taken when certain thresholds are exceeded.

Given the significant impact that resource monitors can have on the operational aspects and billing of a Snowflake account, the capability to create and manage them is restricted to the ACCOUNTADMIN role. This role has the broadest set of privileges in Snowflake, including the ability to manage all aspects of the account, such as users, roles, warehouses, databases, and resource monitors, among others.

References:

Snowflake Documentation on Resource Monitors: [Managing Resource Monitors](#)

Question: 611

What happens to the privileges granted to Snowflake system-defined roles?

- A. The privileges cannot be revoked.
- B. The privileges can be revoked by an ACCOUNTADMIN.
- C. The privileges can be revoked by an orgadmin.
- D. The privileges can be revoked by any user-defined role with appropriate privileges.

Answer: A

Explanation:

The privileges granted to Snowflake's system-defined roles cannot be revoked. System-defined roles, such as SYSADMIN, ACCOUNTADMIN, SECURITYADMIN, and others, come with a set of predefined privileges that are essential for the roles to function correctly within the Snowflake environment.

These privileges are intrinsic to the roles and ensure that users assigned these roles can perform the necessary tasks

and operations relevant to their responsibilities.

The design of Snowflake's role-based access control (RBAC) model ensures that system-defined roles have a specific set of non-revocable privileges to maintain the security, integrity, and operational efficiency of the Snowflake environment. This approach prevents accidental or intentional modification of privileges that could disrupt the essential functions or compromise the security of the Snowflake account.

References:

Snowflake Documentation on Access Control: Understanding Role-Based Access Control (RBAC)

Question: 612

Which type of role can be granted to a share?

- A. Account role
- B. Custom role
- C. Database role
- D. Secondary role

Answer: B

Explanation:

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

Question: 613

When unloading data with the COPY INTO <location> command, what is the purpose of the PARTITION BY <expression> parameter option?

- A. To sort the contents of the output file by the specified expression.
- B. To delimit the records in the output file using the specified expression.
- C. To include a new column in the output using the specified window function expression.
- D. To split the output into multiple files, one for each distinct value of the specified expression.

Answer: D

Explanation:

The PARTITION BY <expression> parameter option in the COPY INTO <location> command is used to split the output into multiple files based on the distinct values of the specified expression. This feature is particularly useful for organizing large datasets into smaller, more manageable files and can help with optimizing downstream processing or consumption of the data. For example, if you are unloading a large dataset of transactions and use PARTITION BY DATE(transactions.transaction_date), Snowflake generates a separate output file for each unique transaction date, facilitating easier data management and access.

This approach to data unloading can significantly improve efficiency when dealing with large volumes of data by enabling parallel processing and simplifying data retrieval based on specific criteria or dimensions.

References:

Snowflake Documentation on Unloading Data: COPY INTO <location>

Question: 614

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

Explanation:

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

Question: 615

Which views are included in the data_sharing_usage schema? (Select TWO).

- A. ACCESS_HISTORY
- B. DATA_TRANSFER_HISTORY
- C. WAREHOUSE_METERING_HISTORY
- D. MONETIZED_USAGE_DAILY
- E. LISTING TELEMETRY DAILY

Answer: DE

Explanation:

<https://docs.snowflake.com/en/sql-reference/data-sharing-usage>

Question: 616

How does the Access_History view enhance overall data governance pertaining to read and write operations? (Select TWO).

- A. Shows how the accessed data was moved from the source to the target objects
- B. Provides a unified picture of what data was accessed and when it was accessed
- C. Protects sensitive data from unauthorized access while allowing authorized users to access it at query runtime
- D. Identifies columns with personal information and tags them so masking policies can be applied to protect sensitive data
- E. Determines whether a given row in a table can be accessed by the user by filtering the data based on a given policy

Answer: B E

Explanation:

The ACCESS_HISTORY view in Snowflake is a powerful tool for enhancing data governance, especially concerning monitoring and auditing data access patterns for both read and write operations. The key ways in which ACCESS_HISTORY enhances overall data governance are:

B . Provides a unified picture of what data was accessed and when it was accessed: This view logs details about query executions, including the objects (tables, views) accessed and the timestamps of these accesses. It's instrumental in auditing and compliance scenarios, where understanding the access patterns to sensitive data is critical.

E . Determines whether a given row in a table can be accessed by the user by filtering the data based on a given policy: While this option is a bit of a misinterpretation of what ACCESS_HISTORY directly offers, it indirectly supports data governance by providing the information necessary to analyze access patterns. This analysis can then inform policy decisions, such as implementing Row-Level Security (RLS) to restrict access to specific rows based on user roles or attributes.

ACCESS_HISTORY does not automatically apply data masking or tag columns with personal information. However, the insights derived from analyzing ACCESS_HISTORY can be used to identify sensitive data and inform the application of masking policies or other security measures.

References:

Snowflake Documentation on ACCESS_HISTORY: Access History

Question: 617

Which Snowflake feature or tool helps troubleshoot issues in SQL query expressions that commonly

cause performance bottlenecks?

- A. Persisted query results
- B. QUERY_HISTORY View
- C. Query acceleration service
- D. Query Profile

Answer: D

Explanation:

The Snowflake feature that helps troubleshoot issues in SQL query expressions and commonly identify performance bottlenecks is the Query Profile. The Query Profile provides a detailed breakdown of a query's execution plan, including each operation's time and resources consumed. It visualizes the steps involved in the query execution, highlighting areas that may be causing inefficiencies, such as full table scans, large joins, or operations that could benefit from optimization. By examining the Query Profile, developers and database administrators can identify and troubleshoot performance issues, optimize query structures, and make informed decisions about potential schema or indexing changes to improve performance.

References:

Snowflake Documentation on Query Profile: Using the Query Profile

Question: 618

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_PRESIGNED_URL
- D. GET_STAGE_LOCATION

Answer: C

Explanation:

The function in Snowflake that returns the URL of a stage using the stage name as the input is C. GET_PRESIGNED_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.

Question: 619

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: B

Explanation:

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

Question: 620

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: C

Explanation:

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

Question: 621

Which Snowflake object contains all the information required to share a database?

- A. Private listing
- B. Secure view
- C. Sequence
- D. Share

Answer: D

Explanation:

In Snowflake, a Share is the object that contains all the information required to share a database with other Snowflake accounts. Shares are used to securely share data stored in Snowflake tables and views, enabling data providers to grant data consumers access to their datasets without duplicating data. When a database is shared, it can include one or more schemas, and each schema can contain tables, views, or both.

References:

Question: 622

Based on Snowflake recommendations, when creating a hierarchy of custom roles, the top-most custom role should be assigned to which role?

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

Answer: A

Explanation:

Based on Snowflake recommendations, when creating a hierarchy of custom roles, the top-most custom role should ideally be granted to the ACCOUNTADMIN role. This recommendation stems from the best practices for implementing a least privilege access control model, ensuring that only the necessary permissions are granted at each level of the role hierarchy. The ACCOUNTADMIN role has the highest level of privileges in Snowflake, including the ability to manage all aspects of the Snowflake account. By assigning the top-most custom role to ACCOUNTADMIN, you ensure that the administration of role hierarchies and the assignment of roles remain under the control of users with the highest level of oversight and responsibility within the Snowflake environment.

References:

Snowflake Documentation on Access Control: Managing Access Control

Question: 623

Which Snowflake table type is only visible to the user who creates it, can have the same name as permanent tables in the same schema, and is dropped at the end of the session?

- A. Temporary
- B. Local
- C. User
- D. Transient

Answer: A

Explanation:

In Snowflake, a Temporary table is a type of table that is only visible to the user who creates it, can have the same name as permanent tables in the same schema, and is automatically dropped at the end of the session in which it was created. Temporary tables are designed for transient data processing needs, where data is needed for the duration of a specific task or session but not beyond. Since they are automatically cleaned up at the end of the session, they help manage storage usage efficiently and ensure that sensitive data is not inadvertently persisted.

References:

Snowflake Documentation on Temporary Tables: Temporary Tables

Question: 624

The property `mins_to_bypass_network_policy` is set at which level?

- A. User
- B. Role
- C. Account
- D. Organization

Answer: C

Explanation:

The property `mins_to_bypass_network_policy` is set at the account level in Snowflake. This setting allows administrators to specify a time frame during which users can bypass network policies that have been set on their account. It is particularly useful in scenarios where temporary access needs to be granted from IPs not covered by the existing network policies. By adjusting this property at the account level, Snowflake administrators can manage and enforce network access controls efficiently across the entire account.

References:

Snowflake Documentation on Network Policies: Network Policies

Question: 625

Which privilege is required on a virtual warehouse to abort any existing executing queries?

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

D. MONITOR

Answer: B

Explanation:

The privilege required on a virtual warehouse to abort any existing executing queries is OPERATE. The OPERATE privilege on a virtual warehouse allows a user to perform operational tasks on the warehouse, including starting, stopping, and restarting the warehouse, as well as aborting running queries. This level of control is essential for managing resource utilization and ensuring that the virtual warehouse operates efficiently.

References:

Snowflake Documentation on Access Control: Access Control Privileges

Question: 626

What is the PRIMARY factor that determines the cost of using a virtual warehouse in Snowflake?

- A. The type of SQL statements executed
- B. The number of tables or databases queried
- C. The amount of data stored in the warehouse
- D. The length of time the compute resources in each cluster run

Answer: D

Explanation:

The primary factor that determines the cost of using a virtual warehouse in Snowflake is D. The length of time the compute resources in each cluster run. Snowflake's pricing model for compute usage is based on the concept of Snowflake credits, which are consumed based on the time virtual warehouses are running and the size of the warehouses. The more compute resources are utilized and the longer they run, the more credits are consumed, which directly impacts the cost.

References:

Snowflake Documentation on Virtual Warehouses: Understanding and Managing Warehouse Credit Usage

Question: 627

What does the Activity area of Snowsight allow users to do? (Select TWO).

- A. Schedule automated data backups.

- B. Explore each step of an executed query.
- C. Monitor queries executed by users in an account.
- D. Create and manage user roles and permissions.
- E. Access Snowflake Marketplace to find and integrate datasets.

Answer: BC

Explanation:

The Activity area of Snowsight, Snowflake's web interface, allows users to perform several important tasks related to query management and performance analysis. Among the options provided, the correct ones are:

- B . Explore each step of an executed query: Snowsight provides detailed insights into query execution, including the ability to explore the execution plan of a query. This helps users understand how a query was processed, identify performance bottlenecks, and optimize query performance.
- C . Monitor queries executed by users in an account: The Activity area enables users to monitor the queries that have been executed by users within the Snowflake account. This includes viewing the history of queries, their execution times, resources consumed, and other relevant metrics.

These features are crucial for effective query performance tuning and ensuring efficient use of Snowflake's resources.

References:

[Snowflake Documentation on Snowsight: Using Snowsight](#)

Question: 628

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: B

Explanation:

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

Question: 629

In addition to performing all the standard steps to share data, which privilege must be granted on each database referenced by a secure view in order to be shared?

- A. READ
- B. REFERENCES
- C. REFERENCE_USAGE
- D. USAGE

Answer: D

Explanation:

In addition to performing all the standard steps to share data, the USAGE privilege must be granted on each database referenced by a secure view in order to be shared. When sharing a database or

specific objects like secure views, the receiving account needs to have the USAGE privilege on the database and schema to access the shared data. This privilege enables the receiving account to access the database and its schemas but does not allow for any DML operations. It's a prerequisite for accessing any objects within the database.

For a secure view to be part of a share, not only does the view itself need to be shared, but the underlying database (and schema, if applicable) must also be accessible to the recipients. Granting USAGE privilege on the database ensures that the receiving account can access the database in a read-only mode to utilize the shared view.

References:

Snowflake Documentation on Shares: Creating and Managing Shares

Question: 630

Which Snowsight feature can be used to perform data manipulations and transformations using a programming

language?

- A. SnowSQL
- B. Dashboards
- C. Python worksheets
- D. Provider Studio

Answer: C

Explanation:

Python worksheets in Snowsight enable users to perform data manipulations and transformations using the Python programming language directly within the Snowflake environment. This feature integrates the power of Python with Snowflake's data warehousing capabilities, allowing for sophisticated data analysis and manipulation.

Introduction to Python Worksheets:

Python worksheets provide an interactive environment to write and execute Python code.

They are designed to facilitate data science and data engineering tasks.

Functionality:

Users can run Python scripts to manipulate data stored in Snowflake.

It allows for leveraging Python's extensive libraries for data analysis, machine learning, and more.

Integration with Snowflake:

Python worksheets run on Snowflake's compute infrastructure, ensuring scalability and performance.

They can access and manipulate Snowflake tables directly, making them a powerful tool for data transformation.

References:

Snowflake Documentation: Snowsight Python Worksheets

Question: 631

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

Explanation:

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

Role: ACCOUNTADMIN:

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.

References:

Snowflake Documentation: Network Policies

Question: 632

What is the MINIMUM size requirement when creating a Snowpark-optimized virtual warehouse?

A. X-Small

B. Small

C. Medium

D. Large

Answer: B

Explanation:

When creating a Snowpark-optimized virtual warehouse in Snowflake, the minimum size requirement is Small. Snowpark is designed to handle data processing workloads efficiently, and the Small size ensures adequate resources for such tasks.

Virtual Warehouse Sizes:

Snowflake offers different sizes for virtual warehouses, ranging from X-Small to 6X-Large.

Each size corresponds to a specific level of compute resources.

Minimum Size Requirement for Snowpark:

A Small virtual warehouse is the minimum size required to optimize performance and resource allocation for Snowpark workloads.

This ensures that the warehouse has sufficient capacity to handle data processing and transformation tasks efficiently.

References:

Snowflake Documentation: Virtual Warehouse Sizes

Question: 633

Awarding a user which privileges on all virtual warehouses is equivalent to granting the user the global MANAGE WAREHOUSES privilege?

- A. MODIFY, MONITOR and OPERATE privileges
- B. ownership and usage privileges
- C. APPLYBUDGET and audit privileges
- D. MANAGE LISTING ADTOTOLfillment and resolve all privileges

Answer: A

Explanation:

Granting a user the MODIFY, MONITOR, and OPERATE privileges on all virtual warehouses in Snowflake is equivalent to granting the global MANAGE WAREHOUSES privilege. These privileges collectively provide comprehensive control over virtual warehouses.

MODIFY Privilege:

Allows users to change the configuration of the virtual warehouse.

Includes resizing, suspending, and resuming the warehouse.

MONITOR Privilege:

Allows users to view the status and usage metrics of the virtual warehouse.

Enables monitoring of performance and workload.

OPERATE Privilege:

Grants the ability to start and stop the virtual warehouse.

Includes pausing and resuming operations as needed.

References:

Snowflake Documentation: Warehouse Privileges

Question: 634

In Snowflake's data security framework, how does column-level security contribute to the protection of sensitive information? (Select TWO).

- A. Implementation of column-level security will optimize query performance.
- B. Column-level security supports encryption of the entire database.
- C. Column-level security ensures that only the table owner can access the data.
- D. Column-level security limits access to specific columns within a table based on user privileges
- E. Column-level security allows the application of a masking policy to a column within a table or view.

Answer: D, E

Explanation:

Column-level security in Snowflake enhances data protection by restricting access and applying masking policies to sensitive data at the column level.

Limiting Access Based on User Privileges:

Column-level security allows administrators to define which users or roles have access to specific columns within a table.

This ensures that sensitive data is only accessible to authorized personnel, thereby reducing the risk of data breaches.

Application of Masking Policies:

Masking policies can be applied to columns to obfuscate sensitive data.

For example, credit card numbers can be masked to show only the last four digits, protecting the full number from being exposed.

References:

Snowflake Documentation: Column-Level Security

Snowflake Documentation: Dynamic Data Masking

Question: 635

How can a user MINIMIZE Continuous Data Protection costs when using large, high-churn, dimension tables?

- A. Create transient tables and periodically copy them to permanent tables.
- B. Create temporary tables and periodically copy them to permanent tables
- C. Create regular tables with extended Time Travel and Fail-safe settings.
- D. Create regular tables with default Time Travel and Fail-safe settings

Answer: A

Explanation:

To minimize Continuous Data Protection (CDP) costs when dealing with large, high-churn dimension tables in Snowflake, using transient tables is a recommended approach.

Transient Tables: These are designed for data that does not require fail-safe protection. They provide the benefit of reducing costs associated with continuous data protection since they do not have the seven-day Fail-safe period that is mandatory for permanent tables.

Periodic Copying to Permanent Tables: By periodically copying data from transient tables to permanent tables, you can achieve a balance between data protection and cost-efficiency. Permanent tables offer the extended data protection features, including Time Travel and Fail-safe, but these features can be applied selectively rather than continuously, reducing the overall CDP COSTS.

References:

Snowflake Documentation on Transient Tables

Snowflake Documentation on Time Travel & Fail-safe

Question: 636

Which function can be used to convert semi-structured data into rows and columns?

- A. TABLE

- B. FLATTEN
- C. PARSE_JSON
- D. JSON_EXTRACT_PATH_TEXT

Answer: B

Explanation:

To convert semi-structured data into rows and columns in Snowflake, the FLATTEN function is utilized.

FLATTEN Function: This function takes semi-structured data (e.g., JSON) and transforms it into a relational table format by breaking down nested structures into individual rows. This process is essential for querying and analyzing semi-structured data using standard SQL operations.

Example Usage:

```
SELECT
  f.value:attribute1 AS attribute1,
  f.value:attribute2 AS attribute2
FROM
  my_table,
  LATERAL FLATTEN(input => my_table.semi_structured_column) f;
```

References:

Snowflake Documentation on FLATTEN

Question: 637

How does Snowflake utilize clustering information to improve query performance?

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

Answer: A

Explanation:

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

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

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

References:

[Snowflake Documentation on Clustering](#)

[Snowflake Documentation on Micro-partition Pruning](#)

Question: 638

How can staged files be removed during data loading once the files have loaded successfully?

- A. Use the DROP command
- B. Use the purge copy option.
- C. Use the FORCE = TRUE parameter
- D. Use the LOAD UNCERTAIN FILES copy option.

Answer: B

Explanation:

To remove staged files during data loading after they have been successfully loaded, the PURGE copy option is used in Snowflake.

PURGE Option: This option automatically deletes files from the stage after they have been successfully copied into the target table.

Usage:

```
FROM @my_stage
```

```
FILE_FORMAT = (type = 'csv')
```

```
PURGE = TRUE;
```

References:

[Snowflake Documentation on COPY INTO](#)

Question: 639

What can be used to process unstructured data?

- A. External tables
- B. The copy into <table> command
- C. External functions
- D. Snowpipe

Answer: C

Explanation:

To process unstructured data in Snowflake, external functions can be used.

External Functions: These allow you to call external services and processing engines from within Snowflake SQL. External functions can be used to handle complex processing tasks that are not natively supported by Snowflake, including those involving unstructured data.

Implementation: You define an external function in Snowflake that points to an external processing service (e.g., AWS Lambda, Google Cloud Functions).

References:

Snowflake Documentation on External Functions

Question: 640

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

Explanation:

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.

References:

Snowflake Documentation: Introduction to Snowpark

Snowflake Documentation: Virtual Warehouses

Question: 641

What is the benefit of using the STRIP_OUTER_ARRAY parameter with the COPY INTO <table> command when loading data from a JSON file into a table?

- A. It flattens multiple arrays into a single array.
- B. It removes the outer array structure and loads separate rows of data
- C. It transforms a pivoted table into an array.
- D. It tokenizes each data string using the defined delimiters.

Answer: B

Explanation:

The `STRIP_OUTER_ARRAY` parameter in the `COPY INTO <table>` command is used when loading data from a JSON file into a table. This parameter removes the outer array structure from the JSON data and loads separate rows of data into the table.

Understanding the `STRIP_OUTER_ARRAY` Parameter:

JSON files often contain data in an array format where multiple records are nested within a single outer array.

The `STRIP_OUTER_ARRAY` parameter helps in simplifying the loading process by removing this outer array, allowing each element within the array to be loaded as a separate row in the target table.

How It Works:

When the `STRIP_OUTER_ARRAY` parameter is set to `TRUE`, Snowflake treats each item within the outer array as an individual record.

This eliminates the need for additional parsing or transformation steps that would otherwise be required to handle nested arrays.

Example Usage:

```
FROM @my_stage/file.json
```

```
FILE_FORMAT = (TYPE = 'JSON' STRIP_OUTER_ARRAY = TRUE);
```

In this example, the JSON file containing an array of objects is loaded into the table `my_table`.

Each object within the array is loaded as a separate row, without the outer array structure.

Benefits:

Simplifies data loading: By removing the outer array, the data is directly loaded into the table without additional manipulation.

Enhances performance: Streamlines the loading process, reducing the complexity and potential errors in handling nested JSON structures.

References:

Snowflake Documentation: `COPY INTO <table>`

Snowflake Documentation: JSON File Format Options

Question: 642

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

Explanation:

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

Question: 643

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

Explanation:

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 semistructured 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.

References:

Snowflake Documentation: TYPEOF

Snowflake Documentation: VARIANT Data Type

Question: 644

Which access control entity in Snowflake can be created as part of a hierarchy within an account?

- A. Securable object
- B. Role
- C. Privilege
- D. User

Answer: B

Explanation:

In Snowflake, a role is an access control entity that can be created as part of a hierarchy within an account. Roles are used to grant and manage privileges in a structured and scalable manner.

Understanding Roles:

Roles are logical entities that group privileges together.

They are used to control access to securable objects like tables, views, warehouses, and more.

Role Hierarchy:

Roles can be organized into a hierarchy, allowing for the inheritance of privileges.

A role higher in the hierarchy (parent role) can grant its privileges to a lower role (child role), simplifying privilege management.

Creating Roles:

Roles can be created using the CREATE ROLE command.

You can define parent-child relationships by granting one role to another.

Example Usage:

```
CREATE ROLE role1;
```

```
CREATE ROLE role2;
```

```
GRANT ROLE role1 TO role2;
```

In this example, role2 inherits the privileges of role1.

Benefits:

Simplifies privilege management: Hierarchies allow for efficient privilege assignment and inheritance.

Enhances security: Roles provide a clear structure for managing access control, ensuring that privileges are granted appropriately.

References:

Snowflake Documentation: Access Control in Snowflake

Snowflake Documentation: Creating and Managing Roles

Question: 645

When an object is created in Snowflake, who owns the object?

- A. The public role
- B. The user's default role
- C. The current active primary role
- D. The owner of the parent schema

Answer: C

Explanation:

In Snowflake, when an object is created, it is owned by the role that is currently active. This active role is the one that is being used to execute the creation command. Ownership implies full control over the object, including the ability to grant and revoke access privileges. This is specified in Snowflake's documentation under the topic of Access Control, which states that "the role in use at the time of object creation becomes the owner of the object."

References:

Snowflake Documentation: Object Ownership

Question: 646

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: B

Explanation:

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.

References:

Snowflake Documentation: Access History

Snowflake Editions: Snowflake Pricing

Question: 647

Which role is responsible for managing the billing and credit data within Snowflake?

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

Answer: A

Explanation:

The ORGADMIN role in Snowflake is responsible for managing organization-level administrative functions, which include managing billing and credit data. This role has the highest level of administrative privileges and can oversee multiple Snowflake accounts within an organization.

References:

Snowflake Documentation: Account and Organization Roles

Question: 648

Which role must be used to create resource monitors?

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

Answer: B

Explanation:

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.

References:

Snowflake Documentation: Resource Monitors

Question: 649

Which command is used to remove files from either external cloud storage or an internal staged

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

Answer: B

Explanation:

The REMOVE command in Snowflake is used to delete files from either external cloud storage locations or internal stages. This command helps manage staged files by removing them when they are no longer needed, which is useful for maintaining organization and managing storage costs.

References:

Snowflake Documentation: REMOVE Command

Question: 650

What persistent data structures are used by the search optimization service to improve the performance of point lookups?

- A. Micro-partitions
- B. Clustering keys
- C. Equality searches
- D. Search access paths

Answer: D

Explanation:

The search optimization service in Snowflake uses persistent data structures known as search access paths to improve the performance of point lookups. These structures enable efficient retrieval of data by reducing the amount of data scanned during queries.

Search Access Paths:

Search access paths are special indexing structures maintained by the search optimization service.

They store metadata about the distribution of data within tables, enabling faster lookups for specific values.

Point Lookups:

Point lookups involve searching for a specific value within a column.

By leveraging search access paths, Snowflake can quickly locate the exact micro-partition containing the value, minimizing the amount of data scanned.

Performance Improvement:

The use of search access paths significantly reduces query execution time for point lookups.

This is especially beneficial for large tables where scanning all micro-partitions would be computationally expensive.

References:

Snowflake Documentation: Search Optimization Service

Snowflake Documentation: Understanding Search Access Paths

Question: 651

What is a characteristic of a tag associated with a masking policy?

- A. A tag can be dropped after a masking policy is assigned
- B. A tag can have only one masking policy for each data type.
- C. A tag can have multiple masking policies for each data type.
- D. A tag can have multiple masking policies with varying data types

Answer: B

Explanation:

In Snowflake, a tag can be associated with only one masking policy for each data type. This means that for a given data type, you can define a single masking policy to be applied when a tag is used. Tags and masking policies are part of Snowflake's data classification and governance features, allowing for data masking based on the context defined by the tags.

References:

Snowflake Documentation: Tag-Based Masking Policies

Question: 652

Which function unloads data from a relational table to JSON?

- A. TRUNC<ID_NUMBER, -6)
- B. TRUNC(ID_NUMBER, 5)
- C. ID_NUMBER*100
- D. TO_CHAR<ID NUMBER)

Answer: D

Explanation:

To unload data from a relational table to JSON format, you can use the TO_CHAR function. This function converts a number to a character string, which can then be serialized into JSON format. While there isn't a direct function specifically named for unloading to JSON, converting the necessary fields to a string representation is a common step in preparing data for JSON serialization.

References:

Question: 653

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: B

Explanation:

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>

Question: 654

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: A B

Explanation:

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., =).

References:

Snowflake Documentation: Search Optimization Service

Question: 655

Which Query Profile operator is considered a DML operator?

- A. ExternalScan
- B. Flatten
- C. Merge
- D. Sort

Answer: C

Explanation:

The Merge operator is considered a DML (Data Manipulation Language) operator in Snowflake. DML operators are used to modify data in tables, and the Merge operator specifically allows for conditional updates, inserts, or deletes in a target table based on the results of a join with a source table.

References:

Snowflake Documentation: Query Profile Overview

Snowflake Documentation: MERGE Statement

Question: 656

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: AD

Explanation:

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

Question: 657

Masking policies are created at what level in Snowflake?

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

Answer: B

Explanation:

Masking policies in Snowflake are created and applied at the column level. These policies are used to obfuscate sensitive data by masking the values in a specific column. Masking policies can enforce data privacy and protection measures by dynamically masking the data based on the role of the user querying the data.

References:

Question: 658

What would cause different results to be returned when running the same query twice?

- A. SAMPLE is used and the seed is set
- B. sample is used and the seed is not set.
- C. Fraction-based sampling is used.
- D. Fixed-size sampling is used.

Answer: B

Explanation:

When using the SAMPLE clause in a query, if the seed is not set, Snowflake will use a different random seed for each execution of the query. This results in different rows being sampled each time, leading to different results. Setting a seed ensures that the same rows are sampled each time the query is run.

References:

Snowflake Documentation: Sampling

Question: 659

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: A

Explanation:

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.

References:

Snowflake Documentation: Network Policy Management

Question: 660

Which Snowflake table is an implicit object layered on a stage, where the stage can be either internal or external?

- A. Directory table
- B. Temporary table
- C. Transient table
- D. A table with a materialized view

Answer: A

Explanation:

A directory table in Snowflake is an implicit object layered on a stage, whether internal or external. It allows users to query the contents of a stage as if it were a table, providing metadata about the files stored in the stage, such as filenames, file sizes, and last modified timestamps.

References:

Snowflake Documentation: Directory Tables

Question: 661

What are type predicates used for?

- A. Extracting data from a variant column
- B. Casting a value in a variant column to a particular data type
- C. Determining if a value in a variant column is a particular data type
- D. Manipulating objects and arrays in a VARIANT column

Answer: C

Explanation:

Type predicates in Snowflake are used to determine if a value in a VARIANT column is of a particular data type. This is useful when working with semi-structured data stored in VARIANT columns, as it allows for data type validation and conditional processing based on the data type.

References:

Snowflake Documentation: Type Predicates

Question: 662

In Snowflake, what allows users to perform recursive queries?

- A. QUALIFY
- B. LATERAL
- C. PIVOT
- D. CONNECT BY

Answer: D

Explanation:

In Snowflake, the CONNECT BY clause allows users to perform recursive queries. Recursive queries are used to process hierarchical or tree-structured data, such as organizational charts or file systems. The CONNECT BY clause is used in conjunction with the START WITH clause to specify the starting point of the hierarchy and the relationship between parent and child rows.

References:

Snowflake Documentation: Hierarchical Queries

Question: 663

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

Explanation:

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.

References:

Snowflake Documentation: RESULT_SCAN

Question: 664

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

Explanation:

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

Question: 665

Which Snowflake native tool can be used to diagnose and troubleshoot network connections?

- A. SnowSQL
- B. Snowflake Python connector
- C. Snowsight
- D. SnowCD

Answer: A

Explanation:

SnowSQL, Snowflake's command-line client, can be used to diagnose and troubleshoot network connections. SnowSQL provides various commands and options to test connectivity, configure network settings, and troubleshoot issues related to network connections between the client and Snowflake.

References:

Snowflake Documentation: SnowSQL

Question: 666

Why would a Snowflake user load JSON data into a VARIANT column instead of a string column?

- A. A VARIANT column is more secure than a string column
- B. A VARIANT column compresses data and a string column does not.
- C. A variant column can be used to create a data hierarchy and a string column cannot
- D. A VARIANT column will have a better query performance than a string column.

Answer: C

Explanation:

A VARIANT column in Snowflake is specifically designed to store semi-structured data, such as JSON, and allows for the creation of a data hierarchy. Unlike string columns, VARIANT columns can natively handle JSON data structures, enabling complex querying and manipulation of hierarchical data using

functions designed for semi-structured data.

References:

Question: 667

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: D, A

Explanation:

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

Question: 668

How can a 5 GB table be downloaded into a single file MOST efficiently?

- A. Keep the default MAX_FILE_SIZE to 16 MB
- B. Set the default MAX_FILE_SIZE to 5 GB.
- C. Set the SINGLE parameter to TRUE.
- D. Use a regular expression in the stage specifications of the COPY command.

Answer: C

Explanation:

To download a 5 GB table into a single file most efficiently in Snowflake, you should set the SINGLE parameter to TRUE. This parameter ensures that the COPY INTO command outputs the result into a single file, regardless of the file size. This approach is more efficient than relying on the default MAX_FILE_SIZE setting, which would split the output into multiple files.

References:

Snowflake Documentation: COPY INTO <location>

Question: 669

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: A, D

Explanation:

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

References:

Snowflake Documentation: Access Control in Snowflake

Question: 670

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

Explanation:

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

Question: 671

What is the MINIMUM Snowflake edition that supports database replication?

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

Answer: B

Explanation:

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

Question: 672

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

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

Answer: B

Explanation:

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

References:

Snowflake Documentation: Query Acceleration Service

Question: 673

Which Snowflake function and command combination should be used to convert rows in a relational table to a single VARIANT column, and unload the rows into a file in JSON format? (Select TWO).

- A. PUT
- B. GET
- C. COPY
- D. EXPORT
- E. OBJECT CONSTRUCT

Answer: C E

Explanation:

To convert rows in a relational table to a single VARIANT column and unload the rows into a file in JSON format, you can use the COPY command in combination with the OBJECT_CONSTRUCT function. The OBJECT_CONSTRUCT function converts the row into a JSON object stored in a VARIANT column, and the COPY command can then be used to unload this data into a JSON file.

References:

Snowflake Documentation: OBJECT_CONSTRUCT

Snowflake Documentation: COPY INTO <location>

Top of Form

Bottom of Form

Question: 674

What Snowflake recommendation is designed to ensure that staged data is only loaded once"?

- A. Partitioning staged data files
- B. Loading only the most recently-staged data files
- C. Removing data files after loading
- D. Identifying and removing duplicates after each data load

Answer: C

Explanation:

Snowflake recommends removing data files from the staging area after they have been loaded into the target table. This practice ensures that the data is only loaded once and prevents accidental reloading of the same data. By removing the files, you eliminate the risk of duplicate data loads.

Stage the Data: Upload the data files to a Snowflake stage (internal or external).

Load the Data: Use the COPY INTO command to load the data from the stage into the Snowflake table.

Remove the Data Files: After successfully loading the data, remove the data files from the stage using the REMOVE command.

References:

Snowflake Documentation: Loading Data into Snowflake

Snowflake Documentation: Staging Data Files

Snowflake Documentation: COPY INTO Command

Question: 675

Which privilege grants the ability to set a column-level security masking policy on a table or view column?

- A. APPLY
- B. CREATE
- C. SET
- D. MODIFY

Answer: A

Explanation:

In Snowflake, the APPLY privilege is required to set a masking policy on a table or view column. This privilege allows a user to associate a masking policy with a specific column, thereby controlling how data in that column is masked or hidden.

Create a Masking Policy: Define a masking policy using the CREATE MASKING POLICY command.

Grant APPLY Privilege: Grant the APPLY privilege on the masking policy to the relevant roles or users.

Apply the Masking Policy: Use the ALTER TABLE or ALTER VIEW command to apply the masking policy to a column.

References:

Snowflake Documentation: Data Masking

Snowflake Documentation: CREATE MASKING POLICY

Snowflake Documentation: Privileges for Masking Policies

Question: 676

How can the Query Profile be used to troubleshoot a problematic query?

- A. It will indicate if a virtual warehouse memory is too small to run the query
- B. It will indicate if a user lacks the privileges needed to run the query.
- C. It will indicate if a virtual warehouse is in auto-scale mode
- D. It will indicate if the user has enough Snowflake credits to run the query

Answer: A

Explanation:

The Query Profile in Snowflake provides detailed insights into the execution of a query. It helps in troubleshooting performance issues by showing the steps of the query execution and the resources consumed. One of the key aspects it reveals is whether the virtual warehouse memory was sufficient for the query.

Access Query Profile: Navigate to the Query History page and select the query you want to analyze.

Examine Query Execution Steps: The Query Profile displays the different stages of the query execution, including the time taken and resources used at each step.

Identify Memory Issues: Look for indicators of memory issues, such as spilling to disk or memory errors, which suggest that the virtual warehouse memory might be too small.

References:

Snowflake Documentation: Query Profile

Snowflake Documentation: Optimizing Queries

Question: 677

Which data type can be used for floating-point numbers without losing precision?

- A. BINARY
- B. VARIANT
- C. INTEGER
- D. DOUBLE

Answer: D

Explanation:

In Snowflake, the DOUBLE data type is used for floating-point numbers and can represent a wide range of values without losing precision. This data type is ideal for storing numerical values that require decimal precision.

Define the Column: When creating a table, specify the column with the DOUBLE data type to store floating-point numbers.

```
CREATE TABLE example_table (  
    id INTEGER,  
    value DOUBLE  
);
```

Insert Data: Insert floating-point numbers into the DOUBLE column.

```
INSERT INTO example_table (id, value) VALUES (1, 123.456);
```

References:

Snowflake Documentation: Data Types

Snowflake Documentation: Numeric Data Types

Question: 678

What does Snowflake recommend as a best practice for using secure views?

- A. Use sequence-generated values
- B. Programmatically reveal the identifiers.
- C. Use secure views solely for query convenience.
- D. Do not expose the sequence-generated column(s)

Answer: D

Explanation:

Snowflake recommends not exposing sequence-generated columns in secure views. Secure views are used to protect sensitive data by ensuring that users can only access data for which they have permissions. Exposing sequence-generated columns can potentially reveal information about the underlying data structure or the number of rows, which might be sensitive.

Create Secure Views: Define secure views using the `SECURE` keyword to ensure they comply with Snowflake's security policies.

Exclude Sensitive Columns: When creating secure views, exclude columns that might expose sensitive information, such as sequence-generated columns.

```
CREATE SECURE VIEW secure_view AS
```

```
SELECT col1, col2
```

```
FROM sensitive_table
```

```
WHERE sensitive_column IS NOT NULL;
```

References:

Snowflake Documentation: Secure Views

Snowflake Documentation: Creating Secure Views

These answers and explanations should provide comprehensive guidance on the specified Snowflake topics.

Question: 679

Which objects will incur storage costs associated with Fail-safe?

- A. External tables
- B. Permanent tables
- C. Data files available in internal stages
- D. Data files available in external stages

Answer: D

Explanation:

Snowflake's Data Exchange allows users to create and manage a group of accounts to which they can offer data shares.

This platform facilitates secure and governed data sharing within and between organizations.

Data Exchange Setup: Users can set up a Data Exchange and invite other Snowflake accounts to join.

Sharing Data: Within the Data Exchange, users can offer shares to the entire group or to specific accounts.

Governance and Security: Data Exchange provides tools for data governance and security, ensuring that shared data is only accessible to authorized accounts.

References:

Snowflake Documentation: Snowflake Data Exchange

Snowflake Documentation: Managing a Data Exchange

Question: 680

Which data sharing option allows a Snowflake user to set up and manage a group of accounts and offer a share to that group?

- A. Free listing
- B. Paid listing
- C. Direct share
- D. Data Exchange

Answer: C

Explanation:

Snowsight, the Snowflake user interface, supports a variety of chart types for data visualization, including line charts and pie charts. These chart types help users analyze and interpret data more effectively.

Line Charts: Useful for showing trends over time or comparing different data sets.

Pie Charts: Ideal for displaying proportions and percentages within a dataset.

References:

Snowflake Documentation: Visualizing Data with Snowsight

Snowflake Documentation: Chart Types

Question: 681

Which type of charts are supported by Snowsight? {Select TWO}

- A. Flowcharts
- B. Gantt charts
- C. Line charts
- D. Pie charts
- E. Scatterplots

Answer: D, E

Explanation:

Question: 682

How can a Snowsight user change a Standard virtual warehouse to a Snowpark-optimized virtual warehouse?

- A. Use the ALTER WAREHOUSE command on an active Standard virtual warehouse
- B. Use the alter warehouse command on an active Snowpark-optimized warehouse.
- C. Use the ALTER warehouse command on a suspended Standard virtual warehouse.
- D. Use the ALTER WAREHOUSE command on a suspended Snowpark-optimized warehouse.

Answer: C

Explanation:

To change a Standard virtual warehouse to a Snowpark-optimized virtual warehouse, the warehouse must be in a suspended state. This ensures that no operations are disrupted during the modification process.

Suspend the Warehouse: Ensure that the warehouse is suspended.

```
ALTER WAREHOUSE my_warehouse SUSPEND;
```

Alter the Warehouse: Modify the warehouse to be Snowpark-optimized.

```
ALTER WAREHOUSE my_warehouse SET WAREHOUSE_TYPE = 'SNOWPARK-OPTIMIZED';
```

Resume the Warehouse: Resume the warehouse to make it operational.

```
ALTER WAREHOUSE my_warehouse RESUME;
```

References:

Snowflake Documentation: Creating Snowpark-Optimized Warehouses

Snowflake Documentation: ALTER WAREHOUSE

Question: 683

What kind of authentication do Snowpipe REST endpoints use?

- A. OAuth
- B. Key-based
- C. Username and password
- D. Single Sign-On (SSO)

Answer: B

Explanation:

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

Question: 684

What are the possible values within a METADATASACTION column in a Snowflake stream? (Select TWO).

A. INSERT

B. UPDATE

- C. DELETE
- D. TRUNCATE
- E. UPSERT

Answer: AC

Explanation:

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

Question: 685

What takes the highest precedence in Snowflake file format options, when specified in multiple locations during data loading?

- A. The stage definition
- B. The table definition
- C. The use of a copy into <table> statement
- D. The use of a copy INTO <location> statement

Answer: C

Explanation:

When loading data into Snowflake, the file format options specified in the COPY INTO <table> statement take the highest precedence over other locations such as the stage or table definitions. This ensures that any specific settings for a particular load operation are applied correctly.

File Format Hierarchy:

Stage Definition: Specifies default file format options for files staged in the location.

Table Definition: Can specify default file format options associated with the table.

COPY INTO Statement: Overrides both the stage and table definitions with the file format options specified directly in

the statement.

Example Usage:

```
COPY INTO my_table
```

```
FROM @my_stage
```

```
FILE_FORMAT = (FORMAT_NAME = 'my_format' FIELD_OPTIONALLY_ENCLOSED_BY = '');
```

References:

Snowflake Documentation: Copy into Table

Snowflake Documentation: File Format Options

Question: 686

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

A. Snowsight

B. SnowCD

C. Snowpark

D. SnowSQL

Answer: D

Explanation:

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

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

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

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

References:

Snowflake Documentation: SnowSQL

Snowflake Documentation: Installing SnowSQL

Question: 687

What Snowflake objects can contain custom application logic written in JavaScript? (Select TWO)

- A. Stored procedures
- B. Stages
- C. Tasks
- D. Views
- E. User-Defined Functions (UDFs)

Answer: A, E

Explanation:

Snowflake allows users to write custom application logic in JavaScript for two types of objects: Stored Procedures and User-Defined Functions (UDFs).

Stored Procedures: Snowflake stored procedures can be written in JavaScript to encapsulate complex business logic and procedural operations.

```
CREATE OR REPLACE PROCEDURE my_procedure()
```

```
RETURNS STRING
```

```
LANGUAGE JAVASCRIPT
```

```
EXECUTE AS CALLER
```

```
AS
```

```
$$
```

```
// JavaScript logic here
```

```
$$;
```

User-Defined Functions (UDFs): Snowflake UDFs can be written in JavaScript to perform custom calculations or operations on data.

```
CREATE OR REPLACE FUNCTION my_function(x FLOAT)
```

```
RETURNS FLOAT
```

```
LANGUAGE JAVASCRIPT
```

```
AS
```

```
$$
```

```
return x * 2;
```

```
$$;
```

References:

Snowflake Documentation: Stored Procedures

Snowflake Documentation: User-Defined Functions (UDFs)

Question: 688

Which Snowflake object can be used to record DML changes made to a table?

- A. Snowpipe
- B. Stage
- C. Stream
- D. Task

Answer: C

Explanation:

Snowflake Streams are used to track and record Data Manipulation Language (DML) changes made to a table. Streams capture changes such as inserts, updates, and deletes, which can then be processed by other Snowflake objects or external applications.

Creating a Stream:

```
CREATE OR REPLACE STREAM my_stream ON TABLE my_table;
```

Using Streams: Streams provide a way to process changes incrementally, making it easier to build efficient data pipelines.

Consuming Stream Data: The captured changes can be consumed using SQL queries or Snowflake tasks.

References:

Snowflake Documentation: Using Streams

Snowflake Documentation: Change Data Capture (CDC) with Streams

Question: 689

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 RSA_P8_KEY='MIIBIjANBgkqh...';
- B. ALTER USER jsmith SET ENCRYPTED_KEY='MIIBIjANBgkqh...';
- C. ALTER USER jsmith SET RSA_PRIVATE_KEY='MIIBIjANBgkqh...';
- D. ALTER USER jsmith SET RSA_PUBLIC_KEY='MIIBIjANBgkqh...';

Answer: D

To use key pair authentication in Snowflake, you need to set the public key for the user. This allows the user to authenticate using their private key.

Generate Key Pair: Generate a public and private key pair.

Set Public Key:

```
ALTER USER jsmith SET RSA_PUBLIC_KEY='MIIBIjANBgkqh...';
```

Authentication: The user can now authenticate by signing requests with the corresponding private key.

References:

Snowflake Documentation: Key Pair Authentication & Key Rotation

Snowflake Documentation: ALTER USER

Question: 690

Secured Data Sharing is allowed for which Snowflake database objects? (Select TWO).

- A. Tables
- B. User-Defined Table Functions (UDTFs)
- C. Secure views
- D. Stored procedures
- E. Worksheets

Answer: AC

Explanation:

Snowflake allows secure data sharing for specific database objects to ensure data is shared securely and efficiently. The primary objects that can be shared securely are tables and secure views.

Tables: Share actual data stored in tables.

Secure Views: Share derived data while protecting the underlying table structures and any sensitive information.

References:

Snowflake Documentation: Introduction to Secure Data Sharing

Snowflake Documentation: Creating Secure Views

Question: 691

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: D, A

Explanation:

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

Question: 692

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

Explanation:

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>

Question: 693

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

Explanation:

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](#)

[Snowflake Documentation: Exporting Data](#)

Question: 694

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

Explanation:

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_PRESIGNED_URL(@my_stage/file.txt);
```

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

References:

Snowflake Documentation: [GET_PRESIGNED_URL](#)

Snowflake Documentation: [Working with Stages](#)

Question: 695

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: BD

Explanation:

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

Question: 696

When unloading data, which combination of parameters should be used to differentiate between empty strings and NULL values? (Select TWO).

- A. ESCAPE_UNENCLOSED_FIELD
- B. REPLACE_INVALID_CHARACTERS
- C. FIELD_OPTIONALLY_ENCLOSED_BY
- D. EMPTY_FIELD_AS_NULL
- E. SKIP_BLANK_LINES

Answer: C, D

Explanation:

When unloading data in Snowflake, it is essential to differentiate between empty strings and NULL values to preserve data integrity. The parameters FIELD_OPTIONALLY_ENCLOSED_BY and EMPTY_FIELD_AS_NULL are used together to address this:

FIELD_OPTIONALLY_ENCLOSED_BY: This parameter specifies the character used to enclose fields, which can differentiate between empty strings (as enclosed fields) and NULLs.

EMPTY_FIELD_AS_NULL: By setting this parameter, Snowflake interprets empty fields as NULL values when unloading data, ensuring accurate representation of NULLs versus empty strings.

These parameters are crucial when exporting data for systems that need explicit differentiation between NULL and empty string values.

Question: 697

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

Explanation:

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.

Question: 698

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

Explanation:

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.

Question: 699

What does the Remote Disk I/O statistic in the Query Profile indicate?

- A. Time spent reading from the result cache.
- B. Time spent reading from the virtual warehouse cache.
- C. Time when the query processing was blocked by remote disk access.
- D. The level of network activity between the Cloud Services layer and the virtual warehouse.

Answer: C

Explanation:

The Remote Disk I/O statistic in the Query Profile reflects time spent waiting on remote disk access, which can occur when data needs to be retrieved from external storage (remote). This metric is crucial for identifying bottlenecks related to I/O delays, often suggesting a need for performance optimization in data retrieval paths.

The other options relate to caching and network activity, but Remote Disk I/O specifically measures the wait time for data access from remote storage locations.

Question: 700

Which type of Snowflake virtual warehouse provides 16 times the memory for each node and is recommended for larger workloads like Machine Learning (ML) training?

- A. A size 6XL warehouse
- B. A standard warehouse
- C. A multi-cluster warehouse
- D. A Snowpark-optimized warehouse

Answer: D

Explanation:

The Snowpark-optimized warehouse is designed with significantly higher memory per node, making it ideal for machine learning (ML) training and other memory-intensive workloads. Snowpark-optimized warehouses provide the necessary computational power and memory for complex, large-scale data processing tasks.

Standard and multi-cluster warehouses do not offer the same memory advantage, while the 6XL warehouse is large but not specifically optimized for memory-intensive operations.

Question: 701

Which Query Profile metrics will provide information that can be used to improve query performance? (Select TWO).

- A. Synchronization
- B. Remote disk IO
- C. Local disk IO
- D. Pruning
- E. Spillage

Answer: B, D

Explanation:

Two key metrics in Snowflake's Query Profile that provide insights for performance improvement are:

Remote Disk IO: This measures the time the query spends waiting on remote disk access, indicating potential performance issues related to I/O bottlenecks.

Pruning: This metric reflects how effectively Snowflake's micro-partition pruning is reducing the data scanned. Better pruning (more partitions excluded) leads to faster query performance, as fewer micro-partitions need to be processed.

These metrics are essential for identifying and addressing inefficiencies in data retrieval and storage access, optimizing overall query performance.

Question: 702

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

Explanation:

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.

Question: 703

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: A, D

Explanation:

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.

Question: 704

A Snowflake table that is loaded using a Kafka connector has a schema consisting of which two variant columns? (Select TWO).

- A. RECORD_TIMESTAMP
- B. RECORD_CONTENT
- C. RECORDKEY
- D. RECORD_SESSION
- E. RECORD_METADATA

Answer: A, C

Explanation:

When using the Snowflake Kafka connector, the table schema includes two important variant columns:

RECORD_TIMESTAMP: This column stores the timestamp from the Kafka record, enabling time-based analysis of incoming data.

RECORDKEY: This captures the unique key of each Kafka message, useful for uniquely identifying records or managing deduplication.

These columns ensure that each message's metadata and key information are preserved, facilitating data analysis and real-time processing tasks in Snowflake.

Question: 705

What best practice recommendations will help prevent timeouts when using the PUT command to load large data sets? (Select TWO).

- A. Compress the files before loading.
- B. Use a semi-structured file format.
- C. Increase the PARALLEL option value.
- D. Load the data into a table stage.
- E. Enable the overwrite option.

Answer: A, C

Explanation:

To avoid timeouts during large data uploads with the PUT command in Snowflake, it is recommended to:

Compress files before loading: Compressed files are smaller and upload faster, reducing the risk of timeouts.

Increase the PARALLEL option value: This option allows more simultaneous upload threads, improving upload speed and efficiency for large datasets.

Semi-structured file formats and table staging do not directly impact timeouts, while enabling overwrite does not prevent timeouts but rather controls overwriting of existing files.

Question: 706

What are characteristics of the ownership privilege when it is granted on a regular Snowflake schema? (Select TWO).

- A. It is automatically granted to the role that creates a database object within the schema.
- B. It allows a role to manage grants on the schema.
- C. It can be transferred from one role to another for a specific schema.
- D. It grants the ability to query data from the schema.
- E. It must be granted to a role in order to alter warehouse settings.

Answer: A, B

Explanation:

In Snowflake, the ownership privilege for a schema includes:

Automatic granting to the creator's role: The role that creates a database object within the schema automatically receives ownership of that object.

Ability to manage grants: Ownership enables the role to manage permissions and grants on the schema and its objects, allowing them to control access at the schema level.

Ownership does not directly confer query privileges or the ability to alter warehouse settings, nor is it transferable without specific privilege management actions.

Question: 707

Who can access the data published in a Data Exchange?

- A. Only the data provider who published the data
- B. Any Snowflake user, regardless of their role or permissions
- C. Any user who has been provided with a unique Data Exchange URL
- D. Only the users that the data provider has invited to the Data Exchange

Answer: D

Explanation:

Snowflake's Data Exchange allows data providers to publish data that is accessible only to users they explicitly invite. This mechanism ensures that only approved users or Snowflake accounts have access, adding a layer of security and control. Invited users can query the shared data directly within their Snowflake account without requiring data copies, reducing redundancy and enhancing data accessibility within controlled boundaries. This selective sharing feature is part of Snowflake's secure data-sharing capabilities.

Question: 708

A Snowflake user accidentally deleted a table. The table no longer exists, but the session is within the data retention period. How can the table be restored using the LEAST amount of operational overhead?

- A. Clone the table schema as it existed before the table was dropped.
- B. Clone the database as it existed before the table was dropped.
- C. Recreate the table and reload the data.
- D. Run the UNDROP command against the table.

Answer: D

Explanation:

In Snowflake, if a table is accidentally dropped but still within the data retention period (also known as "Time Travel"), the simplest and most efficient recovery method is the UNDROP command. This command restores the deleted table to its previous state with minimal operational effort. Since Snowflake retains dropped table data for a specific retention period (up to 90 days for the Enterprise edition), UNDROP can quickly recover the table without the need for complex cloning or data reloading processes, making it ideal for accidental deletions.

Question: 709

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

Explanation:

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.

Question: 710

When used with the UNLOAD command, which parameter specifies the destination of unloaded data?

- A. COPY INTO <table name>
- B. COPY INTO <stage name>
- C. GET <file name>
- D. PUT <file name>

Answer: B

Explanation:

In Snowflake, the COPY INTO <stage name> syntax is used with the UNLOAD command to specify the target location where the data should be unloaded, typically a stage or cloud storage (such as Amazon S3 or Azure Blob Storage). This command unloads data from a Snowflake table into files within the specified destination, enabling easy export and external storage of data. GET and PUT commands are used for file management but are not related to unloading table data directly.

Question: 711

What happens when a table or schema with a standard retention period is dropped?

- A. The object is immediately removed from the system.
- B. The object is instantaneously moved to Fail-safe.
- C. The object is retained but all associated data is immediately purged.
- D. The object is retained for the data retention period.

Answer: D

Explanation:

In Snowflake, when a table or schema is dropped, it is not immediately deleted but retained for the configured data retention period, also known as "Time Travel." During this period, users can use commands like UNDROP to recover the dropped object if needed. After the retention period expires, the object is then moved to Fail-safe (if applicable) for an additional seven days before being permanently removed. This feature is intended to provide data protection and recovery options in case of accidental deletions.

Question: 712

How can a data provider validate that a secure view is configured to display only the data the provider wishes to expose?

- A. Log in to the data consumer account and check if the secure view data is appearing as expected.
- B. Create a data share for a test data consumer account and check if the secure view data is appearing as expected.
- C. Query the secure view from a consumer account by setting the share_restrictions parameter.
- D. Simulate querying the secure view by setting the simulated_data_sharing_consumer session parameter.

Answer: B

Explanation:

The most effective way for a data provider to validate secure view configurations is to create a data share for a test data consumer account. This method allows the provider to review and confirm that only the intended data is accessible in the secure view. Secure views are designed to mask or restrict data visibility, so creating a test share replicates the consumer's experience and ensures data security before sharing with actual consumers.

Question: 713

Which table function will identify data that was loaded using COPY INTO <table> statements and also identify data loaded using Snowpipe?

- A. DATA_TRANSFER_HISTORY
- B. PIPE_USAGE_HISTORY
- C. VALIDATE_PIPE_LOAD
- D. COPY_HISTORY

Answer: D

Explanation:

The COPY_HISTORY table function in Snowflake is used to track data loading activities, including both manual COPY INTO <table> commands and automated data loading through Snowpipe. This function provides visibility into load operations, making it valuable for monitoring data ingestion, ensuring data integrity, and troubleshooting loading processes.

Question: 714

How can the outer array structure of a semi-structured file be removed?

- A. Use the parameter `strip_outer_array = true` in a `COPY INTO <table>` command.
- B. Set the file format to eliminate any outer array structure before initiating the `COPY INTO <table>` command.
- C. Filter the outer array structure using a `PUT` command with the `include_outer_array = false` parameter.
- D. Use the `FLATTEN` command with the `outer_array = false` parameter.

Answer: A

Explanation:

In Snowflake, the parameter `strip_outer_array = true` can be set in the `COPY INTO <table>` command to remove the outer array structure from a semi-structured file. This parameter is useful for JSON files or similar data formats where an array might wrap the data, allowing Snowflake to directly load the inner elements as table rows.

Question: 715

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

Explanation:

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.

Question: 716

Which data protection feature should only be used when all other data recovery options have been attempted?

- A. Time Travel

- B. Cloning
- C. Replication
- D. Fail-safe

Answer: D

Explanation:

The Fail-safe feature in Snowflake is the last-resort data recovery option that should only be used when all other options, such as Time Travel or Cloning, have been exhausted. Fail-safe provides a 7- day recovery period after the Time Travel window has expired but requires Snowflake Support intervention for data recovery. It is intended for disaster recovery and carries additional storage COSTS.

Question: 717

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

Explanation:

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.

Question: 718

How is Single Sign-On (SSO) authentication used in Snowflake?

- A. 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.
- B. SSO is an authentication method that uses a username and password in the API request header.
- C. SSO is an authentication method that allows a user to sign into multiple applications with a single set of credentials.
- D. SSO is an integrated Snowflake feature powered by the Duo Security service which is managed completely by Snowflake.

Answer: C

Explanation:

Question: 719

What action should be taken if a large number of concurrent queries are queued in a virtual warehouse?

- A. Scale-up by resizing the warehouse.

- B. Scale-out with a multi-cluster warehouse.
- C. Disable auto-suspend on the warehouse.
- D. Enable auto-resume on the warehouse.

Answer: A

Explanation:

Question: 720

Which Snowflake object is supported by both database replication and replication groups?

- A. Pipes
- B. Users
- C. Stages
- D. Materialized views

Answer: A

Explanation:

Question: 721

A query is using more credits than expected. The Query Profile shows that a majority of the query execution is spent on remote disk I/O.

How can this be prevented in the future?

- A. Convert the virtual warehouse to Maximized mode.
- B. Increase the size of the virtual warehouse.
- C. Increase the auto suspend time in the virtual warehouse.
- D. Increase the number of clusters in the virtual warehouse.

Answer: B

Explanation:

Question: 722

What object does Snowflake recommend using when planning to unload similarly-formatted data on a regular basis?

- A. Stream
- B. Task
- C. Storage integration
- D. Named file format

Answer: D

Explanation:

Question: 723

A stream object will advance its offset when it is used in which statement?

- A. SELECT
- B. INSERT
- C. CREATE
- D. COPY INTO <location>

Answer: B

Explanation:

Question: 724

What is a fundamental characteristic of Snowflake micro-partitions?

- A. They can be read directly as files.
- B. They serve as an index for Snowflake tables.
- C. They are sized based on Time Travel requirements.
- D. Once established they cannot be changed.

Answer: C

Explanation:

Question: 725

What situation is likely to cause data spillage when a query is run?

- A. When the query contains multiple filters and no data is returned.
- B. When a virtual warehouse runs out of memory while executing the query.
- C. When the number of queries exceeds the max_concurrency_level parameter setting.
- D. When running queries that exceed the statement_timeout_in_secs parameter setting.

Answer: B

Explanation:

Question: 726

Which virtual warehouse auto-suspend configurations will result in a warehouse that runs continually 24 hours a day, 7 days a week? (Select TWO).

- A. -1
- B. 0
- C. 86400
- D. 172800
- E. NULL

Answer: C, D

Explanation:

Question: 727

Which command can be used to unload data into an external named stage in Snowflake?

- A. PUT
- B. CREATE <object>
- C. COPY INTO <table>
- D. COPY INTO <location>

Answer: D

Explanation:

Question: 728

Which Snowflake feature or service is primarily used for managing and monitoring data and user activities?

- A. Snowsight
- B. SnowSQL
- C. Snowflake Marketplace
- D. Streamlit

Answer: A

Explanation:

Question: 729

How does Snowflake recommend defining a clustering key on a high-cardinality column that includes a 15 digit ID numbered column ID_NUMBER?

- A. TRUNC (ID_NUMBER, -6)
- B. TRUNC (ID_NUMBER, 5)
- C. ID_NUMBER*100
- D. TO CHAR(ID_NUMBER)

Answer: A

Explanation:

Question: 730

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

- A. PARSE JSON
- B. FLATTEN

C. OBJECT_CONSTRUCT

D. OBJECT_AGG

Answer: B

Explanation:

Question: 731

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 where operations

C. A column that appears most frequently in group by operations

D. A column that appears most frequently in aggregate operations

E. A column that appears most frequently in join operations

Answer: A, C

Explanation:

Question: 732

Which function, when added to a select statement will return a randomly-selected, specified number of rows from a table?

A. AVERAGE (<num> ROWS)

B. MEDIAN(<num> ROWS)

C. SAMPLE(<num> ROWS)

D. RANDOM(<num> ROWS)

Answer: D

Explanation:

Question: 733

An external stage many_stage contains many directories including one, app_files that contains CSV files

How can all the CSV files from this directory be moved into table my_table without scanning files that are not needed?

A.

```
COPY INTO my_table FROM my_stage PATTERN('*.csv');
```

B.

```
COPY INTO my_table FROM my_stage app_files PATTERN('*.csv');
```

C.

```
COPY INTO my_table FROM my_stage/app_files PATTERN('*.txt');
```

D.

```
LOAD my_table FROM my_stage PATTERN('*.csv');
```

A. Option A

B. Option B

C. Option C

D. Option D

Answer: B

Explanation:

Question: 734

What does Snowflake recommend when planning virtual warehouse usage for a data load?

A. Load the fewest possible number of large files.

B. Dedicate a separate warehouse for loading data.

C. Increase the size of the warehouse used.

D. Use several single-cluster warehouses.

Answer: B

Explanation:

Question: 735

Which transformation techniques are supported for bulk loading data into Snowflake using the copy into <table> command? (Select TWO)

- A. Column grouping
- B. Column omission
- C. Column reordering
- D. Column aggregation
- E. Selection of a limited number of rows

Answer: B, C

Explanation:

Question: 736

Which command can be performed on a Snowflake secure view?

- A. SELECT

B. INSERT

C. UPDATE

D. DELETE

Answer: B

Explanation:

Question: 737

Which features are included in Snowsight? (Select TWO)

A. Worksheet sharing

B. Referencing SnowSQL

C. Exploring the Snowflake Marketplace

D. Changing the Snowflake account cloud provider

E. Downloading query result data larger than 100 MB

Answer: A, C

Explanation:

Question: 738

Which strategy for optimizing virtual warehouse performance should be used to speed the execution of a very large query?

A. Use the query acceleration service.

B. Enable the search optimization service.

C. Convert to a multi-cluster warehouse.

D. Increase the size of the warehouse.

Answer: A

Explanation:

Question: 739

Refer to the exhibit.

Statistics

Sun proff«H	100 00%
Byte* tunned	40.33 MB

> - ■...lUge loannettl from cacfw 25-58%

Byte A'ltten to result	46.91MB
Partitions sunned	10

Partitions total	10
------------------	----

What do these Query Profile metrics indicate about the query micro-partition pruning?

- A. No micro-partitions were pruned during execution.
- B. The query used the local disk cache and pruning was not required.
- C. The size of the data queried was too small to trigger micro-partition pruning.
- D. Since the bytes scanned is less than the bytes written to results some pruning has occurred.

Answer: C

Explanation:

Question: 740

How can network and private connectivity security be managed in Snowflake?

- A. By setting up network policies with IPv4 IP addresses
- B. By putting the Snowflake URL on the allowed list for get method responses
- C. By manually setting up vulnerability patch management policies
- D. By manually setting up an Intrusion Prevention System (IPS) on each account

Answer: A

Explanation:

Question: 741

What actions can be performed by a consumer account on a shared database? (Select TWO)

- A. Cloning a shared table
- B. Modifying the data in a shared table
- C. Using Time Travel on a shared table
- D. Executing the select statement on a shared table
- E. Joining the data from a shared table with another table

Answer: D, E

Explanation:

Question: 742

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

- A. Query Profile
- B. Query history
- C. Visualization chart
- D. Query acceleration service

Answer: A

Explanation:

Question: 743

A Snowflake user wants to design a series of transformations that need to be executed in a specific order on a given schedule.

Which of the snowflake objects should be used?

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

D. Sequences

Answer: B

Question: 744

The customer table in the T1 database is accidentally dropped.

Which privileges are required to restore this table? (Select TWO).

- A. SELECT privilege on the customer table
- B. ownership privilege on the customer table
- C. All privileges on the customer table
- D. All privileges on the T1 database
- E. CREATE TABLE privilege on the T1 database

Answer: D, E

Explanation:

Question: 745

Which object can be used to query the data loading history for the last 365 days for a Snowflake account?

A.

The SK0HIIAtt.AC30UNI_USASE. CaPY^ZSTORX View

B.

The SNOHFLAFX.ACCOUNT JJSAGE.DATA_TRASSFERJiISTORY View

C.

The database I?;FOPHAT:ON_S~HEMA. "OPYJIISTORY table function

D.

The database INFORM? ~- ■ SCHEMA .■■" _7?w;sFEFjn STORY table function

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: A

Explanation:

Question: 746

A Snowflake user needs to optimize the definition of a secure view, but the user cannot see the view.

Which of the LEAST-PRIVILEGED access or role that should be granted to the user to complete this task?

- A. Grant the user the AYSADMIN role.
- B. Grant the user the ownership privilege on the secure view.
- C. Grant the user the imported privileges privilege on the database.
- D. Grant the user the SHOWFLAKE. object viewer database role.

Answer: D

Explanation:

Question: 747

When do Snowflake object owners lose their ability to make grant decisions?

- A. When the object owner has been granted a database role
- B. When the object owner has been granted an account role
- C. When the object is part of a managed access schema
- D. When the object is part of a regular schema

Answer: C

Explanation:

Question: 748

Which table function will return the output of a previously-run command?

- A. FLATTEN
- B. QUERY_HISTORY
- C. TASK_HISTORY
- D. RESULT_SCAN

Answer: C

Explanation:

Question: 749

Which resource monitor setting will cancel all active queries in a virtual warehouse when the threshold is met?

- A. NOTIF
- B. NOTIFY_USERS
- C. SUSPEND
- D. SUSPEND_IMMEDIATE

Answer: D

Explanation:

Question: 750

Which table type is used in the file processing pipeline to process unstructured data in Snowflake?

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

Answer: B

Explanation:

Question: 751

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

Answer: C

Explanation:

Question: 752

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

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

Answer: D, E

Explanation:

Question: 753

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

- A. Standard
- B. Enterprise
- C. Business Critical

D. Virtual Private Snowflake (VPS)

Answer: B

Explanation:

Question: 754

Which JSON paths are considered to be equivalent in Snowflake? (Select TWO).

A.

`src['customer']/EMAIL"`

B.

`src['CUSTOMER']['Email']`

C.

`SRC:Customer.Email`

D.

`src:customer.email`

E.

`SRC:customer.email`

A. Option A

B. Option B

C. Option C

D. Option D

E. Option E

Answer: A

Question: 755

Which multi-duster virtual warehouse setting will help process queued queries as quickly as possible?

- A. An economy scaling policy
- B. A standard scaling policy
- C. Maximized mode
- D. Auto-scale mode

Answer: C, D

Explanation:

Question: 756

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

- A. Alerts

B. External tables

C. Secure views

D. Materialized views

Answer: A

Explanation:

Question: 757

Which default warehouse configuration has the highest precedence whenever a new session is created by a user?

A. Default warehouse for the user

B. Default warehouse in the configuration file of the client utilities

C. Default warehouse specified on a CLI or in drivers/connectors parameters

D. Default warehouse of the role assigned to the user

Answer: D

Explanation:

Question: 758

Which Snowflake objects use storage? (Select TWO)

A. Regular table

B. Regular view

C. Cached query result

D. Materialized view

E. External table

Answer: A, D

Explanation:

Question: 759

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

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

Answer: C

Explanation:

Question: 760

What is the MINIMUM Snowflake edition required to use the query acceleration service?

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

Answer: B

Explanation:

Question: 761

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

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

Answer: B

Explanation:

Question: 762

Which type of URL gives permanent access to files in cloud storage?

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

Answer: D

Explanation:

Question: 763

Where is metadata management handled in Snowflake?

- A. Cloud Services
- B. Compute
- C. Database Storage
- D. Query Processing

Answer: A

Explanation:

Question: 764

When will Snowflake charge credits for the use of the Cloud Services layer?

- A. Credits will be charged whenever the Cloud Services layer is used.
- B. Credits will be charged only when running a Snowflake-provisioned compute warehouse compute_wh.
- C. Credits will be charged when the daily consumption of cloud services resources exceeds Wo of the daily warehouse usage.

D. Credits will be charged only when a virtual warehouse consumes serverless compute services.

Answer: C

Explanation:

Question: 765

What are valid values for the FIELD_OPTIONALLY_ENCLOSED_BY option in the copy into <location> command used during data unloading? (Select TWO).

- A. Single quote character (')
- B. NULL
- C. 'NULL'
- D. NONE
- E. 'NONE'

Answer: A, B

Explanation:

Question: 766

Why is a federated environment used for user authentication in Snowflake?

- A. To enhance data security and privacy
- B. To provide real-time monitoring of user activities
- C. To separate user authentication from user access
- D. To enable direct integration with external databases

Answer: C

Explanation:

Question: 767

A user creates a stage using the following command:

```
CREATE STAGE mystage
```

DIRECTORY = (ENABLE = TRUE)

FILE_FORMAT = myformat;

What will be the outcome?

- A. A stage with a directory table set to automatically refresh will be created.
- B. A stage with a directory table that has metadata that must be manually refreshed will be created.
- 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. The command will fail to run because the name of the directory table is not specified.

Answer: A

Explanation:

Question: 768

Which privilege is required for a user to be able to view a resource monitor?

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

Answer: C

Explanation:

Question: 769

Which features can be used with the Snowflake Standard edition? (Select TWO).

- A. Materialized views
- B. External functions
- C. Multi-cluster virtual warehouses
- D. Fail-safe
- E. Row-level security

Answer: A, C

Explanation:

Question: 770

When cloning a schema, which Snowflake object will not be included in the clone?

- A. An external stage
- B. A named internal stage
- C. A task
- D. A User-Defined Function (UDF)

Answer: A

Explanation:

Question: 771

Which Snowflake data governance feature supports resource usage monitoring?

- A. Data classification
- B. Column lineage
- C. Access history
- D. Object tagging

Answer: D

Explanation:

Question: 772

What is the expiration period for a file URL used to access unstructured data in cloud storage?

- A. The remainder of the session
- B. An unlimited amount of time
- C. The length of time specified in the expiration_time argument
- D. The same length of time as the expiration period for the query results cache

Answer: D

Explanation:

Question: 773

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

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

Answer: C

Explanation:

Question: 774

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 a multi-cluster virtual warehouse
- B. Using the query acceleration service
- C. Using the search optimization service
- D. Disabling auto-suspend

Answer: B

Explanation:

Question: 775

What kind of value does a User-Defined Function (UDF) return? (Select TWO).

- A. Dictionary
- B. List
- C. Object
- D. Scalar
- E. Tabular

Answer: D, E

Explanation:

Question: 776

Which statement will trigger a stream to advance its offset?

- A. DESCRIBE STREAM
- B. ALTER STREAM
- C. DROP STREAM
- D. CREATE OR REPLACE STREAM

Answer: D

Explanation:

Question: 777

Using which object-level parameters will help limit query processing and concurrency slowdowns? (Select TWO).

A.

MULTI_STATEMENT_CCJNT

B.

ENAELE;"E=Y_

C.

STnTEXEXT TIMEOUT IK SECONDS

D.

STATEMENT_QUECED_T7MEOUT_IN_SECONDS

E.

QUERY_A7 ?-HERAT7 ' NJ7-J<_S2ALE_F?."7F

Answer: C, D

Explanation:

Question: 778

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

A. A count of the number of total values in the micro-partition

B. The range of values for each partition in the micro-partition

C. The range of values for each of the rows in the micro-partition

D. The range of values for each of the columns in the micro-partition

E. The number of distinct values for each column in the micro-partition

Explanation:

**Answer: A,
D**