

COF-C02 Dumps

SnowPro Core Certification Exam (COF-C02)

<https://www.certleader.com/COF-C02-dumps.html>



NEW QUESTION 1

- (Topic 1)

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: BD

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

? SnowPro® Core Certification Exam Study Guide

NEW QUESTION 2

- (Topic 1)

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

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

Answer: BCD

Explanation:

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

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

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

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

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

NEW QUESTION 3

- (Topic 1)

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

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

Answer: D

Explanation:

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

References:

? Snowflake Documentation on Secure Data Sharing

? SnowPro™ Core Certification Companion: Hands-on Preparation and Practice

NEW QUESTION 4

- (Topic 1)

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: BD

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

NEW QUESTION 5

- (Topic 1)

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

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

Answer: CE

Explanation:

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

? Continuous Data Protection | Snowflake Documentation¹

? Data Storage Considerations | Snowflake Documentation²

? Snowflake SnowPro Core Certification Study Guide³

? Snowflake Data Cloud Glossary

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

NEW QUESTION 6

- (Topic 1)

Which of the following describes external functions in Snowflake?

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 7

- (Topic 1)

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

- A. True
- B. False

Answer: B

Explanation:

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

NEW QUESTION 8

- (Topic 1)

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

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

Answer: BE

Explanation:

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

References:

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

? Snowflake Documentation on Dynamic Data Masking¹

NEW QUESTION 9

- (Topic 1)

A user has an application that writes a new Tile to a cloud storage location every 5 minutes. What would be the MOST efficient way to get the files into Snowflake?

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

Answer: D

Explanation:

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

References:

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

NEW QUESTION 10

- (Topic 1)

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: CE

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:

- ? SnowPro Core Exam Prep — Answers to Snowflake's LEVEL UP: Backup and Recovery
- ? Snowflake SnowPro Core Certification Exam Questions Set 10

NEW QUESTION 10

- (Topic 1)

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: AE

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

NEW QUESTION 14

- (Topic 1)

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

- A. True
- B. False

Answer: B

Explanation:

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

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

NEW QUESTION 16

- (Topic 1)

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

NEW QUESTION 17

- (Topic 1)

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

NEW QUESTION 19

- (Topic 1)

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: CD

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

NEW QUESTION 24

- (Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 28

- (Topic 1)

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: AC

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.

NEW QUESTION 30

- (Topic 1)

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.

NEW QUESTION 33

- (Topic 1)

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

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

Answer: AC

Explanation:

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

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

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

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

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

NEW QUESTION 38

- (Topic 1)

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

NEW QUESTION 41

- (Topic 1)

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

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

Answer: D

Explanation:

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

References:

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

? Snowflake Documentation on Data Loading

NEW QUESTION 43

- (Topic 1)

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

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

Answer: AB

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 formats1.

? 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 list1.

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

NEW QUESTION 44

- (Topic 1)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 45

- (Topic 1)

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

- A. True
- B. False

Answer: A

Explanation:

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

NEW QUESTION 50

- (Topic 1)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 55

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 59

- (Topic 1)

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

NEW QUESTION 63

- (Topic 1)

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

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

Answer: A

Explanation:

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

References:

- ? Snowflake Documentation on Unloading Data
- ? Snowflake SnowPro Core: Copy Into Command to Unload Rows to Files in Named Stage

NEW QUESTION 68

- (Topic 1)

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: AE

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.

NEW QUESTION 70

- (Topic 1)

What happens when a virtual warehouse is resized?

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

Answer: A

Explanation:

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

References:

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

? Snowflake Documentation on Virtual Warehouses²

NEW QUESTION 73

- (Topic 1)

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

NEW QUESTION 74

- (Topic 1)

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.

NEW QUESTION 79

- (Topic 1)

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: AD

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

NEW QUESTION 81

- (Topic 1)

A user has unloaded data from Snowflake to a stage

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

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

Answer: A

Explanation:

The list command in Snowflake is used to validate and display the list of files in a specified stage. When a user has unloaded data to a stage, running the list @file

stage command will show all the files that have been uploaded to that stage, allowing the user to verify the data that was unloaded.

References:

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

NEW QUESTION 82

- (Topic 1)

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>

NEW QUESTION 85

- (Topic 1)

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: AC

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

NEW QUESTION 88

- (Topic 1)

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

NEW QUESTION 92

- (Topic 1)

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

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

Answer: DE

Explanation:

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

NEW QUESTION 93

- (Topic 1)

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

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

Answer: B

Explanation:

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

References:

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

? Snowflake Documentation on Clustering¹

NEW QUESTION 97

- (Topic 1)

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)

NEW QUESTION 98

- (Topic 1)

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

- A. `SELECT LAST_SUCCESS_LOGIN, LOGIN_NAME FROM ACCOUNT_USAGE.USERS;`
- B. `SELECT EVENT_TIMESTAMP, USER_NAME FROM table(information_schema.login_history_by_user())`
- C. `SELECT EVENT_TIMESTAMP, USER_NAME FROM ACCOUNT_USAGE.ACCESS_HISTORY;`
- D. `SELECT EVENT_TIMESTAMP, USER_NAME FROM ACCOUNT_USAGE.LOGIN_HISTORY;`

Answer: D

Explanation:

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

NEW QUESTION 102

- (Topic 1)

Which Snowflake partner specializes in data catalog solutions?

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

Answer: A

Explanation:

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

References:

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

? Snowflake's official documentation and partner listings

NEW QUESTION 103

- (Topic 1)

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>

NEW QUESTION 106

- (Topic 1)

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 108

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 111

- (Topic 2)

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

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

Answer: BE

Explanation:

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

NEW QUESTION 114

- (Topic 2)

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.

NEW QUESTION 118

- (Topic 2)

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

NEW QUESTION 123

- (Topic 2)

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.

NEW QUESTION 125

- (Topic 2)

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

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

Answer: AE

Explanation:

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

NEW QUESTION 129

- (Topic 2)

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.

NEW QUESTION 134

- (Topic 2)

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

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

Answer: BE

Explanation:

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

NEW QUESTION 139

- (Topic 2)

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.

NEW QUESTION 144

- (Topic 2)

Which of the following describes a Snowflake stored procedure?

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

Answer: D

Explanation:

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

NEW QUESTION 148

- (Topic 2)

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-multicluster.html#:~:text=Multi%2Dcluster%20warehouses%20enable%20you,during%20peak%20and%20off%20hours>.

NEW QUESTION 153

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 154

- (Topic 2)

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

NEW QUESTION 158

- (Topic 2)

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: AB

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

NEW QUESTION 159

- (Topic 2)

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

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

Answer: AD

Explanation:

Best practices for using the ACCOUNTADMIN role include ensuring that all users with this role use Multi-factor Authentication (MFA) for added security.

Additionally, it is recommended to assign the ACCOUNTADMIN role to at least two users to avoid delays in case of password recovery issues, but to as few users as possible to maintain strict control over account-level operations⁴.

NEW QUESTION 161

- (Topic 2)

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: DE

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.

NEW QUESTION 162

- (Topic 2)

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: BC

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

NEW QUESTION 164

- (Topic 2)

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

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

Answer: B

Explanation:

In Snowflake, each size increase in virtual warehouses doubles the number of servers. Therefore, if a size Small virtual warehouse is made up of two servers, a

Large warehouse, which is two sizes larger, would be made up of eight servers (2 servers for Small, 4 for Medium, and 8 for Large)2.
Size specifies the amount of compute resources available per cluster in a warehouse. Snowflake supports the following warehouse sizes:

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

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

NEW QUESTION 166

- (Topic 2)

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

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

Answer: CE

Explanation:

Network policies in Snowflake can be set at the user level and at the account level2.

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

NEW QUESTION 169

- (Topic 2)

Which of the following describes the Snowflake Cloud Services layer?

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

Answer: A

Explanation:

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

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

NEW QUESTION 173

- (Topic 2)

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

NEW QUESTION 178

- (Topic 2)

A user created a new worksheet within the Snowsight UI and wants to share this with teammates. How can this worksheet be shared?

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

Answer: C

Explanation:

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

NEW QUESTION 183

- (Topic 2)

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

NEW QUESTION 186

- (Topic 2)

What type of query benefits the MOST from search optimization?

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

Answer: C

Explanation:

Search optimization in Snowflake is designed to improve the performance of queries that are selective and involve point lookup operations using equality and IN predicates. It is particularly beneficial for queries that access columns with a high number of distinct values¹.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

NEW QUESTION 189

- (Topic 2)

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

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

Answer: ABD

Explanation:

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

NEW QUESTION 190

- (Topic 2)

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 it2.

NEW QUESTION 191

- (Topic 2)

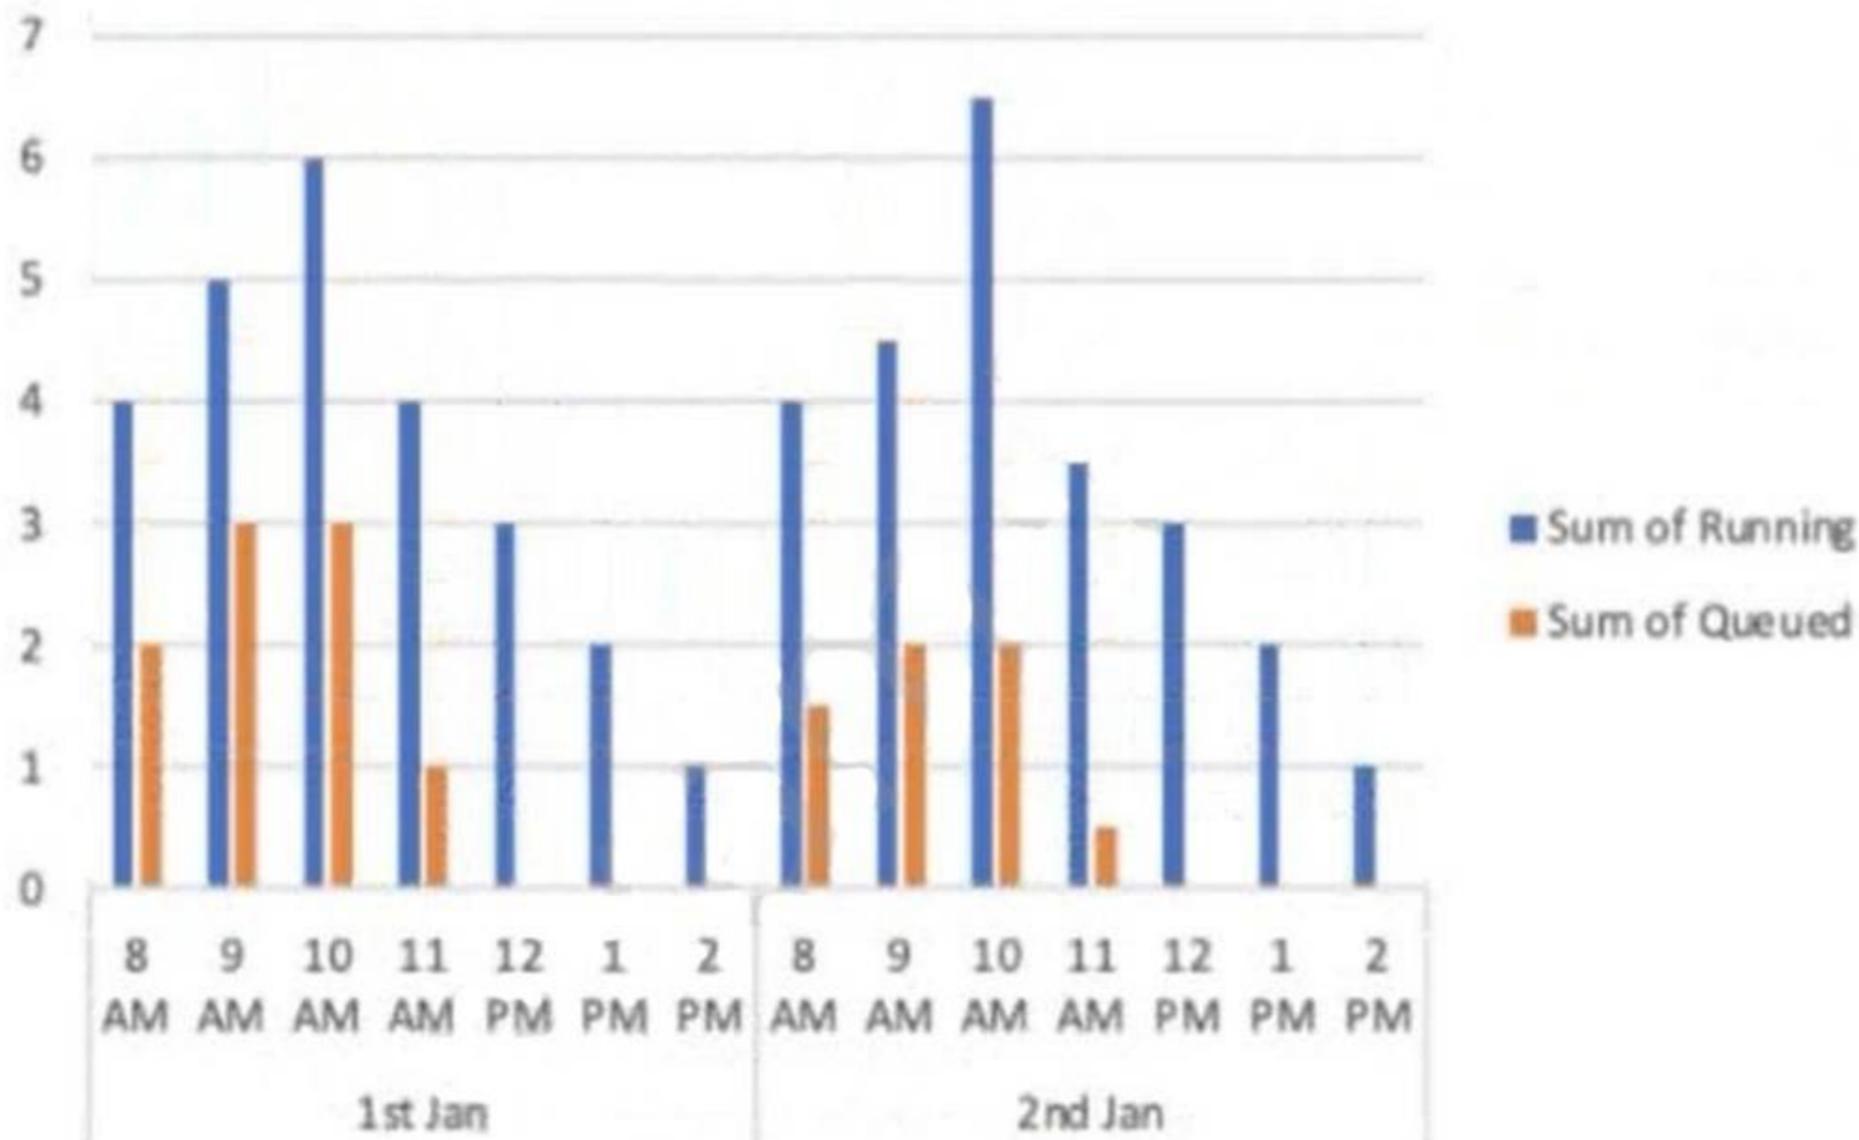
A virtual warehouse is created using the following command:

```
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?

- A. Increase the warehouse size from Medium to 2XL.
- B. Increase the value for the parameter MAX_CONCURRENCY_LEVEL.
- C. Configure the warehouse to a multi-cluster warehouse.
- D. 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

NEW QUESTION 196

- (Topic 2)

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

NEW QUESTION 199

- (Topic 2)

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: AD

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

NEW QUESTION 203

- (Topic 2)

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: BD

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

NEW QUESTION 208

- (Topic 2)

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: AC

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

NEW QUESTION 213

- (Topic 3)

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

- A. It assigns each unloaded data file a unique name.
- B. It uses the put command to download the data by default.
- C. It uses COPY INTO <location> for bulk unloading where the default option is SINGLE - TRUE.
- D. It uses COPY INTO <location> to copy the data from a table into one or more files in an external stage only.

Answer: A

Explanation:

When unloading data, Snowflake assigns each file a unique name to ensure there is no overlap or confusion between files. This is part of the bulk unloading process where data is exported from Snowflake tables into flat files³.

NEW QUESTION 215

- (Topic 3)

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

NEW QUESTION 220

- (Topic 3)

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: CD

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

NEW QUESTION 225

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 228

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 233

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 236

- (Topic 3)

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

NEW QUESTION 241

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 246

- (Topic 3)

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.

NEW QUESTION 250

- (Topic 3)

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

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

Answer: BD

Explanation:

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

NEW QUESTION 254

- (Topic 3)

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

NEW QUESTION 257

- (Topic 3)

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

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

Answer: D

Explanation:

An outer join, specifically a left outer join, will list all rows from the left table and match them with rows from the right table. If there is no match, the result will still

include the row from the left table, with NULLs for columns from the right table. References: Based on general SQL knowledge as of 2021.

NEW QUESTION 262

- (Topic 3)

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: CE

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

NEW QUESTION 265

- (Topic 3)

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: AB

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

NEW QUESTION 270

- (Topic 3)

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: BC

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 searchoptimization service can improve the performance of selective point lookup queries on large tables³⁴.

NEW QUESTION 274

- (Topic 3)

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.

NEW QUESTION 277

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 279

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 284

- (Topic 3)

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 Support.

NEW QUESTION 286

- (Topic 3)

What happens when a database is cloned?

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

Answer: A

Explanation:

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

NEW QUESTION 289

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 293

- (Topic 3)

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 near-real-time data loading, enabling data to be loaded as soon as it's available in the storage layer.

NEW QUESTION 295

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 298

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 300

- (Topic 3)

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

- A. `list @bronzestage;`
- B. `select * from directory(@bronzestage);`
- C. `select metadata$filename from @bronzestage;`
- D. `select * from table(information_schema.stage_directory_file_registration_history(stage name=>'bronzestage1'));`

Answer: B

Explanation:

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

NEW QUESTION 305

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 308

- (Topic 3)

Which of the following are characteristics of security in Snowflake?

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

Answer: C

Explanation:

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

Guide

NEW QUESTION 309

- (Topic 3)

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

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

Answer: AC

Explanation:

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

NEW QUESTION 311

- (Topic 3)

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: AE

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.

NEW QUESTION 314

- (Topic 3)

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

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

Answer: CE

Explanation:

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

NEW QUESTION 318

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 319

- (Topic 3)

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 view. This privilege is necessary before granting the SELECT privilege on a secure view to a share.

NEW QUESTION 320

- (Topic 3)

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.

NEW QUESTION 323

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 328

- (Topic 3)

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: ACF

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 controls.

NEW QUESTION 333

- (Topic 3)

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

How does Snowflake recommend maximizing the load performance?

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

Answer: B

Explanation:

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

[COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 337

- (Topic 3)

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: BE

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.

NEW QUESTION 342

- (Topic 3)

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

NEW QUESTION 345

- (Topic 3)

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

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

Answer: A

Explanation:

Snowsight automatically detects if a target account is in a different region and enables cross-cloud auto-fulfillment when using a paid listing on the Snowflake Marketplace. This feature allows Snowflake to manage the replication of data products to consumer regions as needed, without manual intervention¹.

NEW QUESTION 347

- (Topic 3)

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

NEW QUESTION 352

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 354

- (Topic 3)

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

- A. Table definition
- B. Stage definition

- C. Session level
- D. COPY INTO TABLE statement

Answer: D

Explanation:

When file format options are specified in multiple locations, the load operation applies the options in the following order of precedence: first, the COPY INTO TABLE statement; second, the stage definition; and third, the table definition1

NEW QUESTION 356

- (Topic 3)

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: BC

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

NEW QUESTION 361

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 365

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 369

- (Topic 3)

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 Documentation3.

NEW QUESTION 371

- (Topic 3)

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

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

E. Infrastructure management

Answer: DE

Explanation:

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

NEW QUESTION 372

- (Topic 3)

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

NEW QUESTION 374

- (Topic 3)

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 used. In this case, it would be SHOW GRANTS ON SCHEMA ANALYTICS_DW.MARKETING. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 379

- (Topic 3)

Which of the following statements describes a schema in Snowflake?

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

Answer: A

Explanation:

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

NEW QUESTION 382

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 387

- (Topic 3)

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: BC

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.

NEW QUESTION 389

- (Topic 3)

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.

NEW QUESTION 392

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 393

- (Topic 3)

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.

NEW QUESTION 395

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 399

- (Topic 3)

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

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

D. 16 GB

Answer: B

Explanation:

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

NEW QUESTION 404

- (Topic 3)

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

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

Answer: CD

Explanation:

Snowpark is designed to bring the data programmability to Snowflake, enabling developers to write code in familiar languages like Scala, Java, and Python. It allows for the execution of these codes directly within Snowflake's virtual warehouses, eliminating the need for a separate cluster. Additionally, Snowpark's compatibility with Spark allows users to leverage their existing Spark code with minimal changes¹.

NEW QUESTION 405

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 408

- (Topic 3)

Which of the following describes the Snowflake Cloud Services layer?

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

Answer: A

Explanation:

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

NEW QUESTION 413

- (Topic 3)

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.

NEW QUESTION 416

- (Topic 3)

Which Snowflake URL type is used by directory tables?

- A. File
- B. Pre-signed
- C. Scoped

D. Virtual-hosted style

Answer: C

Explanation:

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

NEW QUESTION 420

- (Topic 3)

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.
- F. An account can be a provider for some datasets and a consumer for others.

Answer: BE

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.

NEW QUESTION 425

- (Topic 3)

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.

NEW QUESTION 426

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 430

- (Topic 3)

What is the name of the SnowSQL file that can store connection information?

- A. history
- B. config
- C. snowsql.cnf
- 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.

NEW QUESTION 434

- (Topic 3)

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

Which condition will trigger a cluster to shut-down?

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

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

Answer: D

Explanation:

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

NEW QUESTION 439

- (Topic 3)

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

NEW QUESTION 441

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 443

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 448

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 450

- (Topic 3)

What is cached during a query on a virtual warehouse?

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

D. All rows accessed during the query

Answer: C

Explanation:

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

NEW QUESTION 455

- (Topic 3)

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 results.

NEW QUESTION 458

- (Topic 3)

Which Snowflake tool would be BEST to troubleshoot network connectivity?

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

Answer: D

Explanation:

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

NEW QUESTION 459

- (Topic 3)

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 Duo.

NEW QUESTION 460

- (Topic 3)

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

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

Answer: BE

Explanation:

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

NEW QUESTION 464

- (Topic 4)

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

- A. COPY INTO

- B. GET
- C. LIST
- D. PUT
- E. REMOVE

Answer: CD

Explanation:

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

NEW QUESTION 465

- (Topic 4)

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

- A. Review the granted privileges to the database object.
- B. Review the row access policy for the database object.
- C. Query the ACCESS_HISTORY view in the ACCOUNT_USAGE schema.
- D. Query the REPLICATION USAGE HISTORY view in the ORGANIZATION USAGE schema.

Answer: C

Explanation:

To determine which user has accessed a database object containing sensitive information, a Snowflake administrator can query the ACCESS_HISTORY view in the ACCOUNT_USAGE schema, which provides information about access to database objects.

NEW QUESTION 468

- (Topic 4)

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

NEW QUESTION 473

- (Topic 4)

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: AE

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

NEW QUESTION 476

- (Topic 4)

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

NEW QUESTION 478

- (Topic 4)

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 filtering². References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 483

- (Topic 4)

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

NEW QUESTION 484

- (Topic 4)

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:

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.

NEW QUESTION 489

- (Topic 4)

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

- A. CREATE TABLE <table_name> LIKE TABLE_A;
- B. CREATE VIEW <view_name> AS SELECT * FROM TABLE_A;
- C. CREATE TABLE <table_name> AS SELECT * FROM TABLE_A;
- D. CREATE MATERIALIZED VIEW <view name> AS SELECT * FROM TABLE A;

Answer: C

Explanation:

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

NEW QUESTION 491

- (Topic 4)

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

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

Answer: AE

Explanation:

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

NEW QUESTION 494

- (Topic 4)

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 load². References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 496

- (Topic 4)

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

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

Answer: AC

Explanation:

The factors that impact storage costs in Snowflake include the account type (Capacity or On Demand) and the cloud region used by the account. These factors determine the rate at which storage is billed, with different regions potentially having different rates³.

NEW QUESTION 499

- (Topic 4)

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

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

Answer: C

Explanation:

The minimum Snowflake Edition that supports secure storage of Protected Health Information (PHI) data is the Business Critical Edition. This edition offers enhanced security features necessary for compliance with regulations such as HIPAA and HITRUST CSF⁴.

NEW QUESTION 502

- (Topic 4)

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.

NEW QUESTION 506

- (Topic 4)

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

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

Answer: AC

Explanation:

Dynamic Data Masking in Snowflake supports tables and views. These objects can have masking policies applied to their columns to dynamically mask data at query time³.

NEW QUESTION 510

- (Topic 4)

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

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

Answer: BD

Explanation:

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

NEW QUESTION 511

- (Topic 4)

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 Travel.

NEW QUESTION 515

- (Topic 4)

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

NEW QUESTION 519

- (Topic 4)

What does the LATERAL modifier for the FLATTEN function do?

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

Answer: C

Explanation:

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

NEW QUESTION 521

- (Topic 4)

Who can grant object privileges in a regular schema?

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

Answer: A

Explanation:

In a regular schema within Snowflake, the object owner has the privilege to grant object privileges. The object owner is typically the role that created the object or to whom the ownership of the object has been transferred. References = [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 522

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your COF-C02 Exam with Our Prep Materials Via below:

<https://www.certleader.com/COF-C02-dumps.html>