

Snowflake

Exam Questions COF-C02

SnowPro Core Certification Exam (COF-C02)



NEW QUESTION 1

- (Topic 1)

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

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

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

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

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

Answer: A

Explanation:

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

References:

? Snowflake Documentation on Data Loading

? SnowPro® Core Certification Study Guide

NEW QUESTION 2

- (Topic 1)

How long is Snowpipe data load history retained?

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

Answer: C

Explanation:

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

References:

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

? Snowflake Documentation on Snowpipe1

NEW QUESTION 3

- (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 4

- (Topic 1)

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

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

Answer: D

Explanation:

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

References:

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

NEW QUESTION 5

- (Topic 1)

What is a limitation of a Materialized View?

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

Answer: D

Explanation:

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

References:

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

NEW QUESTION 6

- (Topic 1)

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 7

- (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 8

- (Topic 1)

Which stage type can be altered and dropped?

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

Answer: B

Explanation:

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

References:

? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Stages1

NEW QUESTION 9

- (Topic 1)

Which of the following describes external functions in Snowflake?

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

Answer: A

Explanation:

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

NEW QUESTION 10

- (Topic 1)

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

- A. True
- B. False

Answer: B

Explanation:

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

NEW QUESTION 10

- (Topic 1)

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

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

Answer: 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 11

- (Topic 1)

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

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

Answer: CE

Explanation:

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

NEW QUESTION 16

- (Topic 1)

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

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

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

Answer: D

Explanation:

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

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

References:

? Snowflake Documentation on Roles

? Snowflake Documentation on Materialized Views

NEW QUESTION 17

- (Topic 1)

What is a machine learning and data science partner within the Snowflake Partner Ecosystem?

- A. Informatica
- B. Power BI
- C. Adobe
- D. Data Robot

Answer: D

Explanation:

Data Robot is recognized as a machine learning and data science partner within the Snowflake Partner Ecosystem. It provides an enterprise AI platform that enables users to build and deploy accurate predictive models quickly. As a partner, Data Robot integrates with Snowflake to enhance data science capabilities.

References:

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

? Snowflake Documentation on Machine Learning & Data Science Partners

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

NEW QUESTION 21

- (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 usage

NEW QUESTION 23

- (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 25

- (Topic 1)

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 30

- (Topic 1)

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

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

Answer: B

Explanation:

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

References:

- ? Snowflake Documentation on Query Results Cache
- ? SnowPro® Core Certification Study Guide

NEW QUESTION 35

- (Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 38

- (Topic 1)

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

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

Answer: DE

Explanation:

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

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Snowflake Documentation on Secure Data Sharing1

NEW QUESTION 40

- (Topic 1)

How often are encryption keys automatically rotated by Snowflake?

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

Answer: A

Explanation:

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

References:

? Understanding Encryption Key Management in Snowflake

NEW QUESTION 43

- (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 44

- (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 47

- (Topic 1)

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

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

Answer: C

Explanation:

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

References:

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

NEW QUESTION 50

- (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 55

- (Topic 1)

Will data cached in a warehouse be lost when the warehouse is resized?

- A. Possibly, if the warehouse is resized to a smaller size and the cache no longer fits.
- B. Yes
- C. because the compute resource is replaced in its entirety with a new compute resource.
- D. No
- E. because the size of the cache is independent from the warehouse size
- F. Yes
- G. because the new compute resource will no longer have access to the cache encryption key

Answer: C

Explanation:

When a Snowflake virtual warehouse is resized, the data cached in the warehouse is not lost. This is because the cache is maintained independently of the warehouse size. Resizing a warehouse, whether scaling up or down, does not affect the cached data, ensuring that query performance is not impacted by such changes. References:

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

? Snowflake Documentation on Virtual Warehouse Performance¹

NEW QUESTION 57

- (Topic 1)

The fail-safe retention period is how many days?

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

Answer: B

Explanation:

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

NEW QUESTION 59

- (Topic 1)

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

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

Answer: D

Explanation:

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

References:

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

? Snowflake Documentation on Snowflake Architecture¹

NEW QUESTION 63

- (Topic 1)

A user is loading JSON documents composed of a huge array containing multiple records into Snowflake. The user enables the strip outer_array file format option. What does the STRIP_OUTER_ARRAY file format do?

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

Answer: B

Explanation:

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

References:

? Snowflake Documentation on JSON File Format

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

NEW QUESTION 66

- (Topic 1)

A sales table FCT_SALES has 100 million records. The following Query was executed

SELECT COUNT (1) FROM FCT SALES;
How did Snowflake fulfill this query?

- A. Query against the result set cache
- B. Query against a virtual warehouse cache
- C. Query against the most-recently created micro-partition
- D. Query against the metadata excite

Answer: D

Explanation:

Snowflake is designed to optimize query performance by utilizing metadata for certain types of queries. When executing a COUNT query, Snowflake can often fulfill the request by accessing metadata about the table's row count, rather than scanning the entire table or micro-partitions. This is particularly efficient for large tables like FCT_SALES with a significant number of records. The metadata layer maintains statistics about the table, including the row count, which enables Snowflake to quickly return the result of a COUNT query without the need to perform a full scan. References:

- ? Snowflake Documentation on Metadata Management
- ? SnowPro® Core Certification Study Guide

NEW QUESTION 71

- (Topic 1)

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

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

Answer: B

Explanation:

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

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Snowflake Documentation on Data Loading1

NEW QUESTION 72

- (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 73

- (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 74

- (Topic 1)

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

- A. True
- B. False

Answer: B

Explanation:

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

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

NEW QUESTION 79

- (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 80

- (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 83

- (Topic 1)

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

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

Answer: DE

Explanation:

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

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

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

? E. The clustering depth for the table is large: A large clustering depth indicates that the table's data is spread across many micro-partitions, which can degrade query performance as more data needs to be scanned.

References:

? Snowflake Documentation on Clustering Keys & Clustered Tables

? Snowflake Documentation on SYSTEM\$CLUSTERING_INFORMATION

? Stack Overflow discussion on cluster key selection in Snowflake

NEW QUESTION 87

- (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 91

- (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 96

- (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.

Parameters#:~:text=The%20parameter%20STATEMENT_QUEUED_TIMEOUT_IN_SECONDS%20sets%20the,indefinitely%20in%20the%20waiting%20queue.

NEW QUESTION 97

- (Topic 1)

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

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

Answer: A

Explanation:

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

References: SHOW GRANTS | Snowflake Documentation

NEW QUESTION 101

- (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 104

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 107

- (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 112

- (Topic 1)

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

- A. True
- B. False

Answer: B

Explanation:

Snowflake does not require source data files to be no larger than 16MB. In fact, Snowflake recommends that for optimal load performance, data files should be roughly 100-250 MB in size when compressed. However, it is not recommended to load very large files (e.g., 100 GB or larger) due to potential delays and wasted credits if errors occur. Smaller files should be aggregated to minimize processing overhead, and larger files should be split to distribute the load among compute resources in an active warehouse. References: Preparing your data files | Snowflake Documentation

NEW QUESTION 115

- (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 118

- (Topic 1)

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

- A. Execute a COPY command
- B. Updating data
- C. Snowpipe
- D. AUTOMATIC CLUSTERING

E. Scaling up a warehouse

Answer: CE

Explanation:

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

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Snowflake Documentation on Snowpipe and Virtual Warehouses1

NEW QUESTION 122

- (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 126

- (Topic 1)

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

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

Answer: B

Explanation:

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

- ? Snowflake Documentation on CREATE MATERIALIZED VIEW1.
- ? Snowflake Documentation on Working with Materialized Views <https://docs.snowflake.com/en/sql-reference/sql/create-materialized-view.html#:~:text=Materialized%20views%20require%20Enterprise%20Edition,upgrading%2C%20please%20contact%20Snowflake%20Support.>

NEW QUESTION 127

- (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 132

- (Topic 1)

When unloading to a stage, which of the following is a recommended practice or approach?

- A. Set SINGLE: = true for larger files
- B. Use OBJECT_CONSTRUCT (*) when using Parquet
- C. Avoid the use of the CAST function
- D. Define an individual file format

Answer: D

Explanation:

When unloading data to a stage, it is recommended to define an individual file format. This ensures that the data is unloaded in a consistent and expected format, which can be crucial for downstream processing and analysis²

NEW QUESTION 137

- (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 140

- (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 144

- (Topic 1)

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

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

Answer: CD

Explanation:

When loading data into Snowflake, it is recommended to:

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

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

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

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

References:

? Snowflake Documentation on Data Loading Considerations

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

NEW QUESTION 147

- (Topic 1)

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

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

Answer: A

Explanation:

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

? SnowPro Core Certification Exam Study Guide

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

NEW QUESTION 152

- (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 154

- (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 158

- (Topic 1)

What is a key feature of Snowflake architecture?

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

Answer: C

Explanation:

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

References:

- ? Snowflake Documentation on Virtual Warehouses
- ? SnowPro® Core Certification Study Guide

NEW QUESTION 159

- (Topic 1)

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

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

Answer: BD

Explanation:

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

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Key Concepts & Architecture | Snowflake Documentation²

NEW QUESTION 164

- (Topic 1)

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

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

Answer: AD

Explanation:

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

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

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

? Snowflake Documentation on Stages

NEW QUESTION 169

- (Topic 2)

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

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

Answer: A

Explanation:

When a pipe is recreated using the CREATE OR REPLACE

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

NEW QUESTION 173

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 175

- (Topic 2)

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

- A. Default Time Travel Retention is set to 0 day
- B. Maximum Time Travel Retention is 30 day
- C. Fail Safe retention time is 1 day.
- D. Default Time Travel Retention is set to 1 day
- E. Maximum Time Travel Retention is 365 day
- F. Fail Safe retention time is 7 days.
- G. Default Time Travel Retention is set to 0 day
- H. Maximum Time Travel Retention is 90 day
- I. Fail Safe retention time is 7 days.
- J. Default Time Travel Retention is set to 1 day
- K. Maximum Time Travel Retention is 90 day
- L. Fail Safe retention time is 7 days.
- M. Default Time Travel Retention is set to 7 day
- N. Maximum Time Travel Retention is 1 day
- O. Fail Safe retention time is 90 days.
- P. Default Time Travel Retention is set to 90 day
- Q. Maximum Time Travel Retention is 7 day
- R. Fail Safe retention time is 356 days.

Answer: D

Explanation:

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

NEW QUESTION 178

- (Topic 2)

Which of the following features are available with the Snowflake Enterprise edition? (Choose two.)

- A. Database replication and failover
- B. Automated index management
- C. Customer managed keys (Tri-secret secure)
- D. Extended time travel
- E. Native support for geospatial data

Answer: AD

Explanation:

The Snowflake Enterprise edition includes database replication and failover for business continuity and disaster recovery, as well as extended time travel capabilities for longer data retention periods1.

NEW QUESTION 179

- (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 180

- (Topic 2)

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

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

Answer: AC

Explanation:

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

NEW QUESTION 185

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 186

- (Topic 2)

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

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

Answer: BC

Explanation:

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

NEW QUESTION 188

- (Topic 2)

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

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

Answer: BC

Explanation:

The Snowflake cloud services layer is responsible for tasks such as authentication and access control, ensuring secure access to the platform, and metadata management, which involves organizing and maintaining information about the data stored in Snowflake56.

NEW QUESTION 189

- (Topic 2)

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

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

Answer: D

Explanation:

The COPY command is used in Snowflake to load data from files located in an external stage into a table. This command allows for efficient and parallelized data loading from various file formats1.

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

NEW QUESTION 193

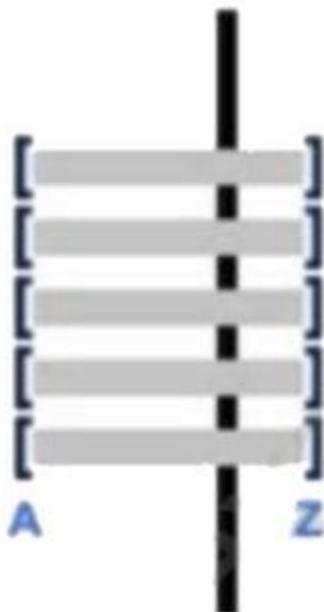
- (Topic 2)

Assume there is a table consisting of five micro-partitions with values ranging from A to Z. Which diagram indicates a well-clustered table?

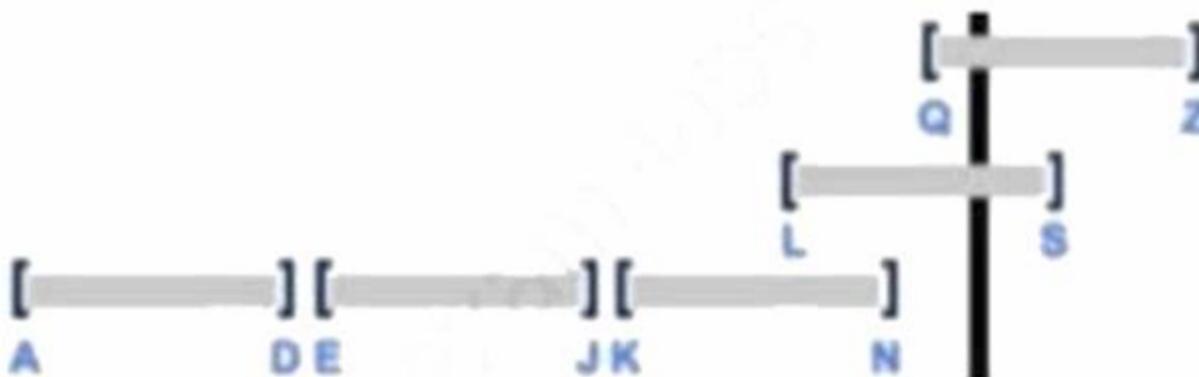
A.



B.



C.



D.



Answer: C

Explanation:

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

References = Snowflake Micro-partitions & Table Clustering

NEW QUESTION 198

- (Topic 2)

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

- A. Role
- B. Stream
- C. Warehouse
- D. External table
- E. User
- F. Share

Answer: BD

Explanation:

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

NEW QUESTION 201

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 202

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 204

- (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 209

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 212

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 215

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 217

- (Topic 2)

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

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

Answer: ACD

Explanation:

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

NEW QUESTION 218

- (Topic 2)

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

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

The load failed with this error:

Max LOB size (16777216) exceeded, actual size of parsed column is 17894470. How can this issue be resolved?

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

Answer: B

Explanation:

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

NEW QUESTION 220

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 224

- (Topic 2)

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

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

Answer: BD

Explanation:

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

NEW QUESTION 225

- (Topic 2)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 227

- (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 230

- (Topic 2)

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

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

Answer: CE

Explanation:

Snowflake virtual warehouses support a configuration file and command line options in SnowSQL to specify a default warehouse, which is characteristic C.

Additionally, the size of a virtual warehouse can be changed at any time, which is characteristic E. These features provide flexibility and ease of use in managing compute resources². References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

NEW QUESTION 234

- (Topic 2)

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

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

Answer: AD

Explanation:

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

NEW QUESTION 239

- (Topic 2)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 240

- (Topic 2)

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

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

Answer: A

Explanation:

When loading data into Snowflake, it is recommended to organize the data into single files with 100-250 MB of compressed data per file. This size range is optimal

for parallel processing and can help in achieving better performance during data loading operations. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 244

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 249

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 253

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 258

- (Topic 2)

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

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

Answer: AD

Explanation:

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

NEW QUESTION 261

- (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 265

- (Topic 2)

Which columns are part of the result set of the Snowflake LATERAL FLATTEN command? (Choose two.)

- A. CONTENT
- B. PATH
- C. BYTE_SIZE
- D. INDEX
- E. DATATYPE

Answer: BD

Explanation:

The LATERAL FLATTEN command in Snowflake produces a result set that includes several columns, among which PATH and INDEX are included. PATH indicates the path to the element within a data structure that needs to be flattened, and INDEX represents the index of the element if it is an array².

NEW QUESTION 268

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 269

- (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 273

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 277

- (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 278

- (Topic 2)

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

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

Answer: DEF

Explanation:

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

NEW QUESTION 281

- (Topic 2)

Which statement is true about running tasks in Snowflake?

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

Answer: B

Explanation:

In Snowflake, a task allows a user to execute a single SQL statement/command using a predefined schedule (B). Tasks are used to automate the execution of SQL statements at scheduled intervals.

NEW QUESTION 285

- (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 286

- (Topic 2)

What impacts the credit consumption of maintaining a materialized view? (Choose two.)

- A. Whether or not it is also a secure view
- B. How often the underlying base table is queried
- C. How often the base table changes
- D. Whether the materialized view has a cluster key defined
- E. How often the materialized view is queried

Answer: CD

Explanation:

The credit consumption for maintaining a materialized view is impacted by how often the base table changes © and whether the materialized view has a cluster key defined (D). Changes to the base table can trigger a refresh of the materialized view, consuming credits. Additionally, having a cluster key defined can optimize the performance and credit usage during the materialized view's maintenance. References: SnowPro Core Certification materialized view credit consumption

NEW QUESTION 287

- (Topic 2)

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

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

Answer: B

Explanation:

The GET command is used to download files from a Snowflake stage to a local folder on a client's machine.
Reference: <https://docs.snowflake.com/en/sql-reference/sql/get.html>

NEW QUESTION 288

- (Topic 2)

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

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

Answer: ABD

Explanation:

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

NEW QUESTION 290

- (Topic 2)

Which of the following accurately describes shares?

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

Answer: A

Explanation:

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

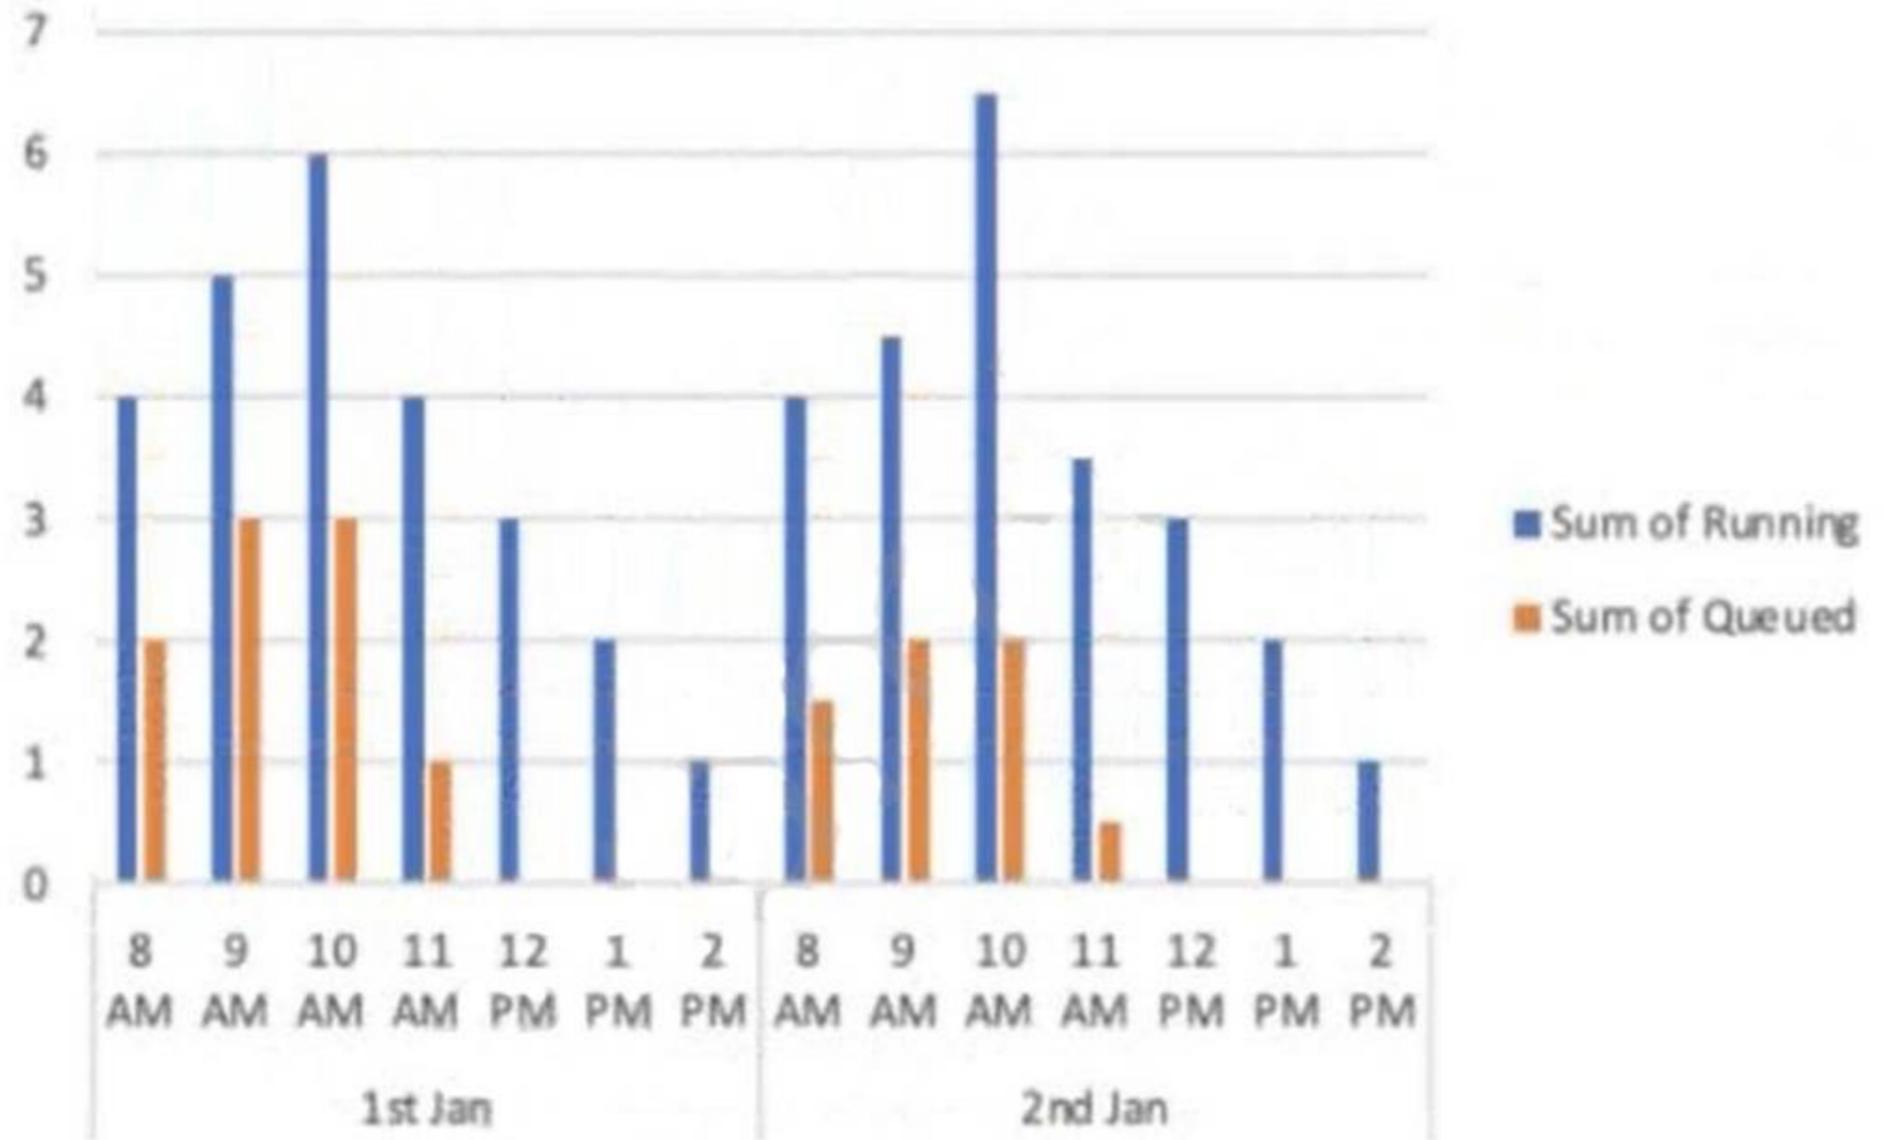
NEW QUESTION 291

- (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 294

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 299

- (Topic 2)

Snowflake supports the use of external stages with which cloud platforms? (Choose three.)

- A. Amazon Web Services
- B. Docker
- C. IBM Cloud
- D. Microsoft Azure Cloud
- E. Google Cloud Platform
- F. Oracle Cloud

Answer: ADE

Explanation:

Snowflake supports the use of external stages with Amazon Web Services (AWS), Microsoft Azure Cloud, and Google Cloud Platform (GCP). These platforms allow users to stage data externally and integrate with Snowflake for data loading operations

NEW QUESTION 300

- (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 301

- (Topic 2)

What is the following SQL command used for? `Select * from table(validate(t1, job_id => '_last'));`

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

Answer: D

Explanation:

The SQL command `Select * from table(validate(t1, job_id => '_last'));` is used to return errors from the last executed COPY command into table t1 in the current session. It checks the results of the most recent data load operation and provides details on any errors that occurred during that process1.

NEW QUESTION 303

- (Topic 2)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 306

- (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 307

- (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 311

- (Topic 2)

Files have been uploaded to a Snowflake internal stage. The files now need to be deleted. Which SQL command should be used to delete the files?

- A. PURGE
- B. MODIFY
- C. REMOVE
- D. DELETE

Answer: C

Explanation:

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

NEW QUESTION 316

- (Topic 2)

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

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

Answer: CE

Explanation:

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

NEW QUESTION 321

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 326

- (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 328

- (Topic 3)

What statistical information in a Query Profile indicates that the query is too large to fit in memory? (Select TWO).

- A. Bytes spilled to local cache.
- B. Bytes spilled to local storage.
- C. Bytes spilled to remote cache.
- D. Bytes spilled to remote storage.

E. Bytes spilled to remote metastore.

Answer: AB

Explanation:

In a Query Profile, the statistical information that indicates a query is too large to fit in memory includes bytes spilled to local cache and bytes spilled to local storage. These metrics suggest that the working data set of the query exceeded the memory available on the warehouse nodes, causing intermediate results to be written to disk

NEW QUESTION 333

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 338

- (Topic 3)

What is used to diagnose and troubleshoot network connections to Snowflake?

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

Answer: A

Explanation:

SnowCD (Snowflake Connectivity Diagnostic Tool) is used to diagnose and troubleshoot network connections to Snowflake. It runs a series of connection checks to evaluate the network connection to Snowflake

NEW QUESTION 343

- (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 346

- (Topic 3)

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

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

Answer: B

Explanation:

The user property `MINS TO BYPASS MFA` in Snowflake allows temporary bypass of MFA for a user, which can be set by an account administrator without contacting Snowflake Support.

NEW QUESTION 349

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 352

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 353

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 355

- (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 359

- (Topic 3)

If a multi-cluster warehouse is using an economy scaling policy, how long will queries wait in the queue before another cluster is started?

- A. 1 minute
- B. 2 minutes
- C. 6 minutes
- D. 8 minutes

Answer: B

Explanation:

In a multi-cluster warehouse with an economy scaling policy, queries will wait in the queue for 2 minutes before another cluster is started. This is to minimize costs by allowing queries to queue up for a short period before adding additional compute resources. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 362

- (Topic 3)

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

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

Answer: BC

Explanation:

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

NEW QUESTION 367

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 369

- (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 search optimization service can improve the performance of selective point lookup queries on large tables³⁴.

NEW QUESTION 374

- (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 376

- (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 381

- (Topic 3)

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

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

Answer: AC

Explanation:

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

NEW QUESTION 384

- (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 Support4.

NEW QUESTION 385

- (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 390

- (Topic 3)

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

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

Answer: B

Explanation:

The query that contains a Snowflake hosted file URL in a directory table for a stage named bronzestage is select * from directory(@bronzestage). This query retrieves a list of all files on the stage along with metadata, including the Snowflake file URL for each file3.

NEW QUESTION 392

- (Topic 3)

Which of the following are characteristics of security in Snowflake?

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

Answer: C

Explanation:

One of the security features of Snowflake includes the periodic rekeying of encrypted data, which is available with the Snowflake Enterprise edition and higher2. This ensures that the encryption keys are rotated regularly to maintain a high level of security. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 396

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 397

- (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 duplication²³. Clones are independent of the original table and do not stay in sync with it, nor do they inherently have better query performance. However, they do inherit the time travel history from the original table at the time of cloning

NEW QUESTION 400

- (Topic 3)

By definition, a secure view is exposed only to users with what privilege?

- A. IMPORT SHARE
- B. OWNERSHIP
- C. REFERENCES
- D. USAGE

Answer: B

Explanation:

A secure view in Snowflake is exposed only to users with the OWNERSHIP privilege. This privilege ensures that only authorized users who own the view, or roles that include ownership, can access the secure view

NEW QUESTION 401

- (Topic 3)

What is a characteristic of the Snowflake Query Profile?

- A. It can provide statistics on a maximum number of 100 queries per week.
- B. It provides a graphic representation of the main components of the query processing.
- C. It provides detailed statistics about which queries are using the greatest number of compute resources.
- D. It can be used by third-party software using the Query Profile API.

Answer: B

Explanation:

The Snowflake Query Profile provides a graphic representation of the main components of the query processing. This visual aid helps users understand the execution details and performance characteristics of their queries⁴.

NEW QUESTION 403

- (Topic 3)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 404

- (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 view². This privilege is necessary before granting the SELECT privilege on a secure view to a share.

NEW QUESTION 409

- (Topic 3)

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

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

Answer: AE

Explanation:

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

NEW QUESTION 414

- (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 417

- (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 controls⁴.

NEW QUESTION 419

- (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 efficiently¹. References:

[COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 421

- (Topic 3)

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

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

Answer: DE

Explanation:

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

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

NEW QUESTION 423

- (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 426

- (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 427

- (Topic 3)

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

- A. Include the SQL statements in the body of the task `CREATE TASK mytask .. AS INSERT INTO target1 SELECT .. FROM stream_s1 WHERE .. INSERT INTO target2 SELECT .. FROM stream_s1 WHERE ..`
- B. A stored procedure can have only one DML statement per stored procedure invocation and therefore the user should sequence stored procedure calls in the taskdefinition `CREATE TASK mytask AS call stored_proc1(); call stored_proc2();`
- C. Use a stored procedure executing multiple SQL statements and invoke the stored procedure from the task
- D. `CREATE TASK mytask AS call stored_proc_multiple_statements_inside();`
- E. Create a task for each SQL statement (e.
- F. resulting in task1, task2, etc.) and string the series of SQL statements by having a control task calling task1, task2, et
- G. sequentially.

Answer: C

Explanation:

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

NEW QUESTION 431

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 433

- (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 437

- (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 438

- (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 441

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 446

- (Topic 3)

How do Snowflake data providers share data that resides in different databases?

- A. External tables
- B. Secure views
- C. Materialized views
- D. User-Defined Functions (UDFs)

Answer: B

Explanation:

Snowflake data providers can share data residing in different databases through secure views. Secure views allow for the referencing of objects such as schemas, tables, and other views contained in one or more databases, as long as those databases belong to the same account. This enables providers to share data

securely and efficiently with consumers. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 447

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 450

- (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 455

- (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 460

- (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 461

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

COF-C02 Practice Exam Features:

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

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