

Snowflake

Exam Questions COF-C02

SnowPro Core Certification Exam (COF-C02)



NEW QUESTION 1

- (Topic 1)

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

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

Answer: BD

Explanation:

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

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

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

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

References:

? Snowflake Documentation on Security Features

? SnowPro® Core Certification Exam Study Guide

NEW QUESTION 2

- (Topic 1)

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

- A. UTF-8
- B. UTF-16
- C. ISO S859-1
- D. ANSI_X3.A

Answer: A

Explanation:

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

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

NEW QUESTION 3

- (Topic 1)

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

- A. Cluster keys
- B. Time Travel
- C. Fail-safe
- D. Cloning

Answer: B

Explanation:

Snowflake's Time Travel feature is used for both querying historical data in tables and restoring and cloning historical data in databases, schemas, and tables. It allows users to access historical data within a defined period (1 day by default, up to 90 days for Snowflake Enterprise Edition) and is a key feature for data recovery and management. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 4

- (Topic 1)

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

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

Answer: BD

Explanation:

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

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

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

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

? Snowflake Documentation on Virtual Warehouses

? SnowPro® Core Certification Study Guide

NEW QUESTION 5

- (Topic 1)

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

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

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

Answer: D

Explanation:

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

References:

? Snowflake Documentation on Snowpipe

? SnowPro® Core Certification Study Guide

NEW QUESTION 6

- (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 capabilities2.

References:

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

? Snowflake Documentation on Machine Learning & Data Science Partners

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

NEW QUESTION 7

- (Topic 1)

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

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

Answer: AE

Explanation:

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

References:

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

? Snowflake's official documentation on billing and usage1

NEW QUESTION 8

- (Topic 1)

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

- A. True
- B. False

Answer: B

Explanation:

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

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

NEW QUESTION 9

- (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 10

- (Topic 1)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 10

- (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 14

- (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 16

- (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 19

- (Topic 1)

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

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

Answer: B

Explanation:

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

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

? Snowflake Documentation on Account Usage Schema¹

NEW QUESTION 22

- (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 23

- (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 27

- (Topic 1)

What is a best practice after creating a custom role?

- A. Create the custom role using the SYSADMIN role.
- B. Assign the custom role to the SYSADMIN role
- C. Assign the custom role to the PUBLIC role
- D. Add CUSTOM to all custom role names

Answer: B

Explanation:

Assigning the custom role to the SYSADMIN role is considered a best practice because it allows the SYSADMIN role to manage objects created by the custom role. This is important for maintaining proper access control and ensuring that the SYSADMIN can perform necessary administrative tasks on objects created by users with the custom role.

References:

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

? Section 1.3 - SnowPro Core Certification Study Guide¹

NEW QUESTION 29

- (Topic 1)

A user unloaded a Snowflake table called mytable to an internal stage called mystage. Which command can be used to view the list of files that has been uploaded to the staged?

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

Answer: D

Explanation:

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

References:

? Snowflake Documentation on Stages

? SnowPro® Core Certification Study Guide

NEW QUESTION 30

- (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 35

- (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. Ye
- C. because the compute resource is replaced in its entirety with a new compute resource.
- D. N
- E. because the size of the cache is independent from the warehouse size
- F. Ye
- G. became 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 40

- (Topic 1)

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

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

Answer: D

Explanation:

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

References:

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

? Snowflake Documentation on Masking Policies

NEW QUESTION 43

- (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 47

- (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 51

- (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 54

- (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 57

- (Topic 1)

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

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

Answer: A

Explanation:

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

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

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

NEW QUESTION 59

- (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 63

- (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 66

- (Topic 1)

When reviewing the load for a warehouse using the load monitoring chart, the chart indicates that a high volume of Queries are always queuing in the warehouse. According to recommended best practice, what should be done to reduce the Queue volume? (Select TWO).

- A. Use multi-clustered warehousing to scale out warehouse capacity.
- B. Scale up the warehouse size to allow Queries to execute faster.
- C. Stop and start the warehouse to clear the queued queries
- D. Migrate some queries to a new warehouse to reduce load
- E. Limit user access to the warehouse so fewer queries are run against it.

Answer: AB

Explanation:

To address a high volume of queries queuing in a warehouse, Snowflake recommends two best practices:

? A. Use multi-clustered warehousing to scale out warehouse capacity: This approach allows for the distribution of queries across multiple clusters within a warehouse, effectively managing the load and reducing the queue volume.

? B. Scale up the warehouse size to allow Queries to execute faster: Increasing the size of the warehouse provides more compute resources, which can reduce the time it takes for queries to execute and thus decrease the number of queries waiting in the queue.

These strategies help to optimize the performance of the warehouse by ensuring that resources are scaled appropriately to meet demand.

References:

? Snowflake Documentation on Multi-Cluster Warehousing

? SnowPro Core Certification best practices

NEW QUESTION 69

- (Topic 1)

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

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

Answer: A

Explanation:

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

References:

? Snowflake Documentation on Pipes

? SnowPro® Core Certification Study Guide <https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro.html>

NEW QUESTION 71

- (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 72

- (Topic 1)

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

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

Answer: D

Explanation:

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

References:

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

? Snowflake Documentation on Query Profile1

? Snowpro Core Certification Exam Flashcards2

NEW QUESTION 77

- (Topic 1)

What happens when a virtual warehouse is resized?

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

Answer: A

Explanation:

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

References:

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

? Snowflake Documentation on Virtual Warehouses2

NEW QUESTION 81

- (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 82

- (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 87

- (Topic 1)

Which statement about billing applies to Snowflake credits?

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

Answer: D

Explanation:

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

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

NEW QUESTION 88

- (Topic 1)

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

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

Answer: DE

Explanation:

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

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

NEW QUESTION 89

- (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 92

- (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_NAMEFROM 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 97

- (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 98

- (Topic 1)

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

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

Answer: A

Explanation:

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

References:

? Snowflake Documentation on Reader Accounts

? SnowPro® Core Certification Study Guide

NEW QUESTION 103

- (Topic 1)

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

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

Answer: A

Explanation:

The ACCOUNTADMIN role is required for a user to create and manage accounts in Snowflake. This role has the highest level of privileges and is responsible for managing all aspects of the Snowflake account, including the ability to create and manage other user accounts1.
<https://docs.snowflake.com/en/user-guide/security-access-control-considerations.html>

NEW QUESTION 104

- (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 pipe2.

NEW QUESTION 108

- (Topic 2)

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

```
{
  "customer": [
    {
      "address": "San Francisco, CA",
      "name": "Jane Doe"
    }
  ],
  "date": "2022-01-28",
  "dealership": "Town Auto Sales",
  "salesperson": {
    "id": "55"
  }
}
```

A user needs to extract the dealership information from the JSON. How can this be accomplished?

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

Answer: B

Explanation:

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

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

NEW QUESTION 109

- (Topic 2)

When publishing a Snowflake Data Marketplace listing into a remote region what should be taken into consideration? (Choose two.)

- A. There is no need to have a Snowflake account in the target region, a share will be created for each user.
- B. The listing is replicated into all selected regions automatically, the data is not.
- C. The user must have the ORGADMIN role available in at least one account to link accounts for replication.
- D. Shares attached to listings in remote regions can be viewed from any account in an organization.
- E. For a standard listing the user can wait until the first customer requests the data before replicating it to the target region.

Answer: BC

Explanation:

When publishing a Snowflake Data Marketplace listing into a remote region, it's important to note that while the listing is replicated into all selected regions automatically, the data itself is not. Therefore, the data must be replicated separately. Additionally, the user must have the ORGADMIN role in at least one account to manage the replication of accounts¹.

NEW QUESTION 110

- (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 111

- (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 114

- (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 117

- (Topic 2)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 118

- (Topic 2)

What affects whether the query results cache can be used?

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

Answer: C

Explanation:

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

NEW QUESTION 119

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 120

- (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 124

- (Topic 2)

What is the purpose of multi-cluster virtual warehouses?

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

Answer: C

Explanation:

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

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

NEW QUESTION 128

- (Topic 2)

What types of data listings are available in the Snowflake Data Marketplace? (Choose two.)

- A. Reader
- B. Consumer
- C. Vendor
- D. Standard
- E. Personalized

Answer: CE

Explanation:

In the Snowflake Data Marketplace, the types of data listings available include ??Vendor??, which refers to the providers of data, and ??Personalized??, which indicates customized data offerings tailored to specific consumer needs⁴⁵.

NEW QUESTION 129

- (Topic 2)

A company needs to allow some users to see Personally Identifiable Information (PII) while limiting other users from seeing the full value of the PII. Which Snowflake feature will support this?

- A. Row access policies

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

Answer: B

Explanation:

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

NEW QUESTION 131

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 133

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 138

- (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 stage¹².
References = DROP STAGE, REMOVE

NEW QUESTION 139

- (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 144

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

NEW QUESTION 147

- (Topic 2)

How are serverless features billed?

- A. Per second multiplied by an automatic sizing for the job
- B. Per minute multiplied by an automatic sizing for the job, with a minimum of one minute
- C. Per second multiplied by the size, as determined by the SERVERLESS_FEATURES_SIZE account parameter
- D. Serverless features are not billed, unless the total cost for the month exceeds 10% of the warehouse credits, on the account

Answer: B

Explanation:

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

NEW QUESTION 149

- (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 152

- (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 156

- (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 credits2.

NEW QUESTION 157

- (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 159

- (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 160

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 163

- (Topic 2)

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

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 166

- (Topic 2)

Which of the following describes the Snowflake Cloud Services layer?

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

Answer: A

Explanation:

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

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

NEW QUESTION 168

- (Topic 2)

How many days is load history for Snowpipe retained?

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

Answer: C

Explanation:

Snowpipe retains load history for 14 days. This allows users to view and audit the data that has been loaded into Snowflake using Snowpipe within this time frame³.

NEW QUESTION 169

- (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 172

- (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 174

- (Topic 2)

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

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

Answer: B

Explanation:

Column-level security, which allows for the application of masking policies to columns in tables or views, is available starting from the Enterprise edition of Snowflake. References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

NEW QUESTION 178

- (Topic 2)

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

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

Answer: AD

Explanation:

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

[COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 179

- (Topic 2)

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

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

Answer: CE

Explanation:

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

References = Introduction to Secure Data Sharing | Snowflake Documentation

NEW QUESTION 181

- (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 186

- (Topic 2)

A user created a transient table and made several changes to it over the course of several days. Three days after the table was created, the user would like to go back to the first version of the table.

How can this be accomplished?

- A. Use Time Travel, as long as DATA_RETENTION_TIME_IN_DAYS was set to at least 3 days.
- B. The transient table version cannot be retrieved after 24 hours.
- C. Contact Snowflake Support to have the data retrieved from Fail-safe storage.
- D. Use the FAIL_SAFE parameter for Time Travel to retrieve the data from Fail-safe storage.

Answer: A

Explanation:

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

NEW QUESTION 187

- (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 192

- (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 193

- (Topic 2)

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 196

- (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 199

- (Topic 2)

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

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

Answer: AE

Explanation:

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

NEW QUESTION 203

- (Topic 2)

Which Snowflake function will interpret an input string as a JSON document, and produce a VARIANT value?

- A. `parse_json()`
- B. `json_extract_path_text()`
- C. `object_construct()`
- D. `flatten`

Answer: A

Explanation:

The `parse_json()` function in Snowflake interprets an input string as a JSON document and produces a VARIANT value containing the JSON document. This function is specifically designed for parsing strings that contain valid JSON information³.

NEW QUESTION 204

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 206

- (Topic 3)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 211

- (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 212

- (Topic 3)

What is a responsibility of Snowflake??s virtual warehouses?

- A. Infrastructure management
- B. Metadata management
- C. Query execution
- D. Query parsing and optimization
- E. Permanent storage of micro-partitions

Answer: C

Explanation:

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

NEW QUESTION 216

- (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 219

- (Topic 3)

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

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

Answer: BCE

Explanation:

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

NEW QUESTION 221

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 223

- (Topic 3)

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

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

Answer: C

Explanation:

The pre-signed URL type allows users or applications to download or access files directly from a Snowflake stage without authentication. This URL type is open and can be used without needing to authenticate into Snowflake or pass an authorization token.

NEW QUESTION 226

- (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 Support².

NEW QUESTION 227

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 229

- (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 233

- (Topic 3)

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

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

Answer: BC

NEW QUESTION 235

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 240

- (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 244

- (Topic 3)

Which statement MOST accurately describes clustering in Snowflake?

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

Answer: B

Explanation:

Clustering in Snowflake refers to the organization of data within micro- partitions, which are contiguous units of storage within Snowflake tables. Clustering keys can be defined to co-locate similar rows in the same micro-partitions, improving scan efficiency and query performance¹². References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 246

- (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 251

- (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 252

- (Topic 3)

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

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

Answer: CE

Explanation:

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

NEW QUESTION 255

- (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 257

- (Topic 3)

Which command is used to unload files from an internal or external stage to a local file system?

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

Answer: B

Explanation:

The command used to unload files from an internal or external stage to a local file system in Snowflake is the GET command. This command allows users to download data files that have been staged, making them available on the local file system for further use²³.

NEW QUESTION 258

- (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 260

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 261

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 263

- (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 267

- (Topic 3)

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

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

Answer: B

Explanation:

The Snowflake feature that allows for small volumes of data to be continuously loaded into Snowflake and incrementally made available for analysis is Snowpipe. Snowpipe is designed for near-real-time data loading, enabling data to be loaded as soon as it's available in the storage layer.

NEW QUESTION 272

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 275

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 276

- (Topic 3)

Which of the following are characteristics of security in Snowflake?

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

Answer: C

Explanation:

One of the security features of Snowflake includes the periodic rekeying of encrypted data, which is available with the Snowflake Enterprise edition and 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 281

- (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 286

- (Topic 3)

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

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

Answer: B

Explanation:

The OPERATE privilege is required for a role to resume a suspended warehouse if auto-resume is not enabled. This privilege allows the role to start, stop, suspend, or resume a virtual warehouse³.

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

NEW QUESTION 287

- (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 289

- (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 293

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 297

- (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 301

- (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 303

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 306

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 308

- (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 311

- (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 315

- (Topic 3)

Which of the following statements describes a schema in Snowflake?

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

Answer: A

Explanation:

A schema in Snowflake is a logical grouping of database objects, such as tables and views, that belongs to a single database. Each schema is part of a namespace in Snowflake, which is inferred from the current database and schema in use for the session⁵

NEW QUESTION 319

- (Topic 3)

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

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

Answer: D

Explanation:

In the Activity section in Snowsight, Snowflake users can explore the performance of executed queries. This includes monitoring queries, viewing details about queries, including performance data, and exploring each step of an executed query in the query profile¹.

NEW QUESTION 320

- (Topic 3)

What action can a user take to address query concurrency issues?

- A. Enable the query acceleration service.
- B. Enable the search optimization service.
- C. Add additional clusters to the virtual warehouse
- D. Resize the virtual warehouse to a larger instance size.

Answer: C

Explanation:

To address query concurrency issues, a user can add additional clusters to the virtual warehouse. This allows for the distribution of queries across multiple clusters, reducing the load on any single cluster and improving overall query performance².

NEW QUESTION 322

- (Topic 3)

Which kind of Snowflake table stores file-level metadata for each file in a stage?

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

Answer: A

Explanation:

The kind of Snowflake table that stores file-level metadata for each file in a stage is a directory table. A directory table is an implicit object layered on a stage and stores file-level metadata about the data files in the stage.

NEW QUESTION 323

- (Topic 3)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 324

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 325

- (Topic 3)

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

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

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

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

Answer: B

Explanation:

The reader account compute will be billed to the provider account.

Very Comprehensive Explanation

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

NEW QUESTION 328

- (Topic 3)

Which of the following describes the Snowflake Cloud Services layer?

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

Answer: A

Explanation:

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

NEW QUESTION 331

- (Topic 3)

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

- A. Scale up the virtual warehouse.

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

Answer: A

Explanation:

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

NEW QUESTION 333

- (Topic 3)

What role is required to use Partner Connect?

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

Answer: A

Explanation:

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

NEW QUESTION 338

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 341

- (Topic 3)

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

- A. PARSE_JSON
- B. CHECK_JSON
- C. JSON_EXTRACT_PATH_TEXT
- D. PARSE_XML

Answer: A

Explanation:

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

NEW QUESTION 342

- (Topic 3)

Which pages are included in the Activity area of Snowsight? (Select TWO).

- A. Contacts
- B. Sharing settings
- C. Copy History
- D. Query History
- E. Automatic Clustering History

Answer: DE

Explanation:

The Activity area of Snowsight includes the Query History page, which allows users to monitor and view details about queries executed in their account, including performance data1. It also includes the Automatic Clustering History, which provides insights into the automatic clustering operations performed on tables2.

NEW QUESTION 346

- (Topic 3)

A view is defined on a permanent table. A temporary table with the same name is created in the same schema as the referenced table. What will the query from the view return?

- A. The data from the permanent table.

- B. The data from the temporary table.
- C. An error stating that the view could not be compiled.
- D. An error stating that the referenced object could not be uniquely identified.

Answer: A

Explanation:

When a view is defined on a permanent table, and a temporary table with the same name is created in the same schema, the query from the view will return the data from the permanent table. Temporary tables are session-specific and do not affect the data returned by views defined on permanent tables2.

NEW QUESTION 351

- (Topic 3)

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

- A. history
- B. config
- C. snowsqlcnf
- D. snowsql.pubkey

Answer: B

Explanation:

The SnowSQL file that can store connection information is named ??config??. It is used to store user credentials and connection details for easy access to Snowflake instances. References: Based on general database knowledge as of 2021.

NEW QUESTION 355

- (Topic 3)

Which commands should be used to grant the privilege allowing a role to select data from all current tables and any tables that will be created later in a schema? (Choose two.)

- A. grant USAGE on all tables in schema DB1.SCHEMA to role MYROLE;
- B. grant USAGE on future tables in schema DB1.SCHEMA to role MYROLE;
- C. grant SELECT on all tables in schema DB1.SCHEMA to role MYROLE;
- D. grant SELECT on future tables in schema DB1.SCHEMA to role MYROLE;
- E. grant SELECT on all tables in database DB1 to role MYROLE;
- F. grant SELECT on future tables in database DB1 to role MYROLE;

Answer: CD

Explanation:

To grant a role the privilege to select data from all current and future tables in a schema, two separate commands are needed. The first command grants the SELECT privilege on all existing tables within the schema, and the second command grants the SELECT privilege on all tables that will be created in the future within the same schema.

NEW QUESTION 358

- (Topic 3)

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

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

Answer: BD

Explanation:

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

NEW QUESTION 362

- (Topic 3)

Credit charges for Snowflake virtual warehouses are calculated based on which of the following considerations? (Choose two.)

- A. The number of queries executed
- B. The number of active users assigned to the warehouse
- C. The size of the virtual warehouse
- D. The length of time the warehouse is running
- E. The duration of the queries that are executed

Answer: CD

Explanation:

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

NEW QUESTION 367

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 369

- (Topic 3)

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

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

Answer: A

Explanation:

Snowflake automatically rotates the Account and Table master keys when they are more than 30 days old. Active keys are retired, and new keys are created, ensuring robust security through frequent key changes¹

NEW QUESTION 373

- (Topic 3)

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

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

Answer: A

Explanation:

To find out which numeric values are present in both tables, the INTERSECT set operator should be used. This operator returns rows from one query's result set which also appear in another query's result set, effectively finding the common elements between the two tables⁴⁵.

NEW QUESTION 374

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 378

- (Topic 3)

Which feature is integrated to support Multi-Factor Authentication (MFA) at Snowflake?

- A. Authy
- B. Duo Security
- C. One Login
- D. RSA SecurID Access

Answer: B

Explanation:

Snowflake integrates Duo Security to support Multi-Factor Authentication (MFA). This feature provides increased login security for users connecting to Snowflake, and it is managed completely by Snowflake without the need for users to sign up separately with Duo⁴.

NEW QUESTION 380

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 385

- (Topic 3)

What technique does Snowflake use to limit the number of micro-partitions scanned by each query?

- A. B-tree
- B. Indexing
- C. Map reduce
- D. Pruning

Answer: D

Explanation:

Snowflake uses a technique called pruning to limit the number of micro-partitions scanned by each query. Pruning effectively filters out unnecessary micro-partitions based on the query's filter conditions, which can significantly improve query performance by reducing the amount of data scanned¹.

NEW QUESTION 386

- (Topic 4)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 388

- (Topic 4)

Which data types can be used in Snowflake to store semi-structured data? (Select TWO)

- A. ARRAY
- B. BLOB
- C. CLOB
- D. JSON
- E. VARIANT

Answer: AE

Explanation:

Snowflake supports the storage of semi-structured data using the ARRAY and VARIANT data types. The ARRAY data type can directly contain VARIANT, and thus indirectly contain any other data type, including itself. The VARIANT data type can store a value of any other type, including OBJECT and ARRAY, and is often used to represent semi-structured data formats like JSON, Avro, ORC, Parquet, or XML³⁴.

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

NEW QUESTION 389

- (Topic 4)

What information is found within the Statistic output in the Query Profile Overview?

- A. Operator tree
- B. Table pruning
- C. Most expensive nodes
- D. Nodes by execution time

Answer: C

Explanation:

The Statistic output in the Query Profile Overview of Snowflake provides detailed insights into the performance of different parts of the query. Specifically, it highlights the "Most expensive nodes," which are the operations or steps within the query execution that consume the most resources, such as CPU and memory. Identifying these nodes helps in pinpointing performance bottlenecks and optimizing query execution by focusing efforts on the most resource-intensive parts of the

query.

References:

? Snowflake Documentation on Query Profile Overview: It details the components of the profile overview, emphasizing how to interpret the statistics section to improve query performance by understanding which nodes are most resource-intensive.

NEW QUESTION 393

- (Topic 4)

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

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

Answer: B

Explanation:

The query acceleration service in Snowflake is designed to benefit queries that involve large scans and selective filters. This service can offload portions of the query processing work to shared compute resources, which can handle these types of workloads more efficiently by performing more work in parallel and reducing the wall-clock time spent in scanning and filtering². References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 397

- (Topic 4)

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

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

Answer: B

Explanation:

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

NEW QUESTION 401

- (Topic 4)

Which Snowflake command can be used to unload the result of a query to a single file?

- A. Use COPY INTO <external stage> followed by a GET command to download the file.
- B. Use COPY INTO <internal stage> followed by a put command to download the file.
- C. Use COPY INTO <internal stage> with SINGLE = TRUE followed by a GET command to download the file.
- D. Use COPY INTO <external stage> with SINGLE = TRUE followed by a PUT command to download the file.

Answer: C

Explanation:

The Snowflake command to unload the result of a query to a single file is COPY INTO <internal stage> with SINGLE = TRUE followed by a GET command to download the file. This command unloads the query result into a single file in the specified internal stage

NEW QUESTION 402

- (Topic 4)

For which use cases is running a virtual warehouse required? (Select TWO).

- A. When creating a table
- B. When loading data into a table
- C. When unloading data from a table
- D. When executing a show command
- E. When executing a list command

Answer: BC

Explanation:

Running a virtual warehouse is required when loading data into a table and when unloading data from a table because these operations require compute resources that are provided by the virtual warehouse²³.

NEW QUESTION 404

- (Topic 4)

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

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

Answer: AC

Explanation:

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

NEW QUESTION 407

- (Topic 4)

Which Snowflake function will parse a JSON-null into a SQL-null?

- A. TO_CHAR
- B. TO_VARIANT
- C. TO_VARCHAR
- D. STRIP NULL VALUE

Answer: D

Explanation:

The STRIP_NULL_VALUE function in Snowflake is used to convert a JSON null value into a SQL NULL value¹.

NEW QUESTION 410

- (Topic 4)

How can a Snowflake user validate data that is unloaded using the COPY INTO <location> command?

- A. Load the data into a CSV file.
- B. Load the data into a relational table.
- C. Use the VALIDATION_MODE - SQL statement.
- D. Use the validation mode = return rows statement.

Answer: C

Explanation:

To validate data unloaded using the COPY INTO <location> command, a Snowflake user can use the VALIDATION_MODE parameter within the SQL statement to test the files for errors without loading them³.

NEW QUESTION 415

- (Topic 4)

What is the primary purpose of a directory table in Snowflake?

- A. To store actual data from external stages
- B. To automatically expire file URLs for security
- C. To manage user privileges and access control
- D. To store file-level metadata about data files in a stage

Answer: D

Explanation:

A directory table in Snowflake is used to store file-level metadata about the data files in a stage. It is conceptually similar to an external table and provides information such as file size, last modified timestamp, and file URL. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 420

- (Topic 4)

What feature of Snowflake Continuous Data Protection can be used for maintenance of historical data?

- A. Access control
- B. Fail-safe
- C. Network policies
- D. Time Travel

Answer: D

Explanation:

Snowflake's Time Travel feature is used for the maintenance of historical data, allowing users to access and restore data that has been changed or deleted within a defined period⁴.

NEW QUESTION 423

- (Topic 4)

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

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

Answer: AC

Explanation:

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

NEW QUESTION 428

- (Topic 4)

A Snowflake account has activated federated authentication.

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

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

Answer: A

Explanation:

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

NEW QUESTION 430

- (Topic 4)

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

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

Answer: BD

Explanation:

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

NEW QUESTION 435

- (Topic 4)

Which parameter can be set at the account level to set the minimum number of days for which Snowflake retains historical data in Time Travel?

- A. DATA_RETENTION_TIME_IN_DAYS
- B. MAX_DATA_EXTENSION_TIME_IN_DAYS
- C. MIN_DATA_RETENTION_TIME_IN_DAYS
- D. MAX_CONCURRENCY_LEVEL

Answer: A

Explanation:

The parameter DATA_RETENTION_TIME_IN_DAYS can be set at the account level to define the minimum number of days Snowflake retains historical data for Time Travel.

NEW QUESTION 438

- (Topic 4)

What is the relationship between a Query Profile and a virtual warehouse?

- A. A Query Profile can help users right-size virtual warehouses.
- B. A Query Profile defines the hardware specifications of the virtual warehouse.
- C. A Query Profile can help determine the number of virtual warehouses available.
- D. A Query Profile automatically scales the virtual warehouse based on the query complexity.

Answer: A

Explanation:

A Query Profile provides detailed execution information for a query, which can be used to analyze the performance and behavior of queries. This information can help users optimize and right-size their virtual warehouses for better efficiency. References:
[COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 439

- (Topic 4)

What is a directory table in Snowflake?

- A. A separate database object that is used to store file-level metadata
- B. An object layered on a stage that is used to store file-level metadata
- C. A database object with grantable privileges for unstructured data tasks
- D. A Snowflake table specifically designed for storing unstructured files

Answer: B

Explanation:

A directory table in Snowflake is an object layered on a stage that is used to store file-level metadata. It is not a separate database object but is conceptually similar to an external table because it stores metadata about the data files in the stage⁵.

NEW QUESTION 442

- (Topic 4)

What does the LATERAL modifier for the FLATTEN function do?

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

Answer: C

Explanation:

The LATERAL modifier for the FLATTEN function allows joining information outside the object (such as other columns in the source table) with the flattened data, creating a lateral view that correlates with the preceding tables in the FROM clause²³⁴⁵. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 444

- (Topic 4)

Which command is used to start configuring Snowflake for Single Sign-On (SSO)?

- A. CREATE SESSION POLICY
- B. CREATE NETWORK RULE
- C. CREATE SECURITY INTEGRATION
- D. CREATE PASSWORD POLICY

Answer: C

Explanation:

To start configuring Snowflake for Single Sign-On (SSO), the CREATE SECURITY INTEGRATION command is used. This command sets up a security integration object in Snowflake, which is necessary for enabling SSO with external identity providers using SAML 2.0¹. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 449

- (Topic 4)

Which Snowflake object does not consume any storage costs?

- A. Secure view
- B. Materialized view
- C. Temporary table
- D. Transient table

Answer: C

Explanation:

Temporary tables do not consume any storage costs in Snowflake because they only exist for the duration of the session that created them and are automatically dropped when the session ends, thus incurring no long-term storage charges⁴. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 451

- (Topic 4)

Which statements describe benefits of Snowflake's separation of compute and storage? (Select TWO).

- A. The separation allows independent scaling of computing resources.
- B. The separation ensures consistent data encryption across all virtual data warehouses.
- C. The separation supports automatic conversion of semi-structured data into structured data for advanced data analysis.
- D. Storage volume growth and compute usage growth can be tightly coupled.
- E. Compute can be scaled up or down without the requirement to add more storage.

Answer: AE

Explanation:

Snowflake's architecture allows for the independent scaling of compute resources, meaning you can increase or decrease the computational power as needed without affecting storage. This separation also means that storage can grow independently of compute usage, allowing for more flexible and cost-effective data management.

NEW QUESTION 455

- (Topic 4)

What happens when a Snowflake user changes the data retention period at the schema level?

- A. All child objects will retain data for the new retention period.
- B. All child objects that do not have an explicit retention period will automatically inherit the new retention period.
- C. All child objects with an explicit retention period will be overridden with the new retention period.
- D. All explicit child object retention periods will remain unchanged.

Answer: B

Explanation:

When the data retention period is changed at the schema level, all child objects that do not have an explicit retention period set will inherit the new retention period from the schema4.

NEW QUESTION 458

.....

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](#)