

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)

What is a limitation of a Materialized View?

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

Answer: D

Explanation:

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

References:

? Snowflake Documentation on Materialized Views

? SnowPro® Core Certification Study Guide

NEW QUESTION 4

- (Topic 1)

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

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

Answer: BD

Explanation:

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

NEW QUESTION 5

- (Topic 1)

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

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

Answer: CE

Explanation:

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

? Continuous Data Protection | Snowflake Documentation¹

? Data Storage Considerations | Snowflake Documentation²

? Snowflake SnowPro Core Certification Study Guide³

? Snowflake Data Cloud Glossary

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

NEW QUESTION 6

- (Topic 1)

What features does Snowflake Time Travel enable?

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

Answer: BC

Explanation:

Snowflake Time Travel is a powerful feature that allows users to access historical data within a defined period. It enables two key capabilities:

? B. Restoring data-related objects that have been deleted within the past 90 days:

Time Travel can be used to restore tables, schemas, and databases that have been accidentally or intentionally deleted within the Time Travel retention period.

? C. Conducting point-in-time analysis for BI reporting: It allows users to query

historical data as it appeared at a specific point in time within the Time Travel retention period, which is crucial for business intelligence and reporting purposes.

While Time Travel does allow querying of past data, it is limited to the retention period set for the Snowflake account, which is typically 1 day for standard accounts

and can be extended up to 90 days for enterprise accounts. It does not enable querying or restoring objects created or deleted beyond the retention period, nor

does it provide analysis over all periods of time.

References:

? Snowflake Documentation on Time Travel

? SnowPro® Core Certification Study Guide

NEW QUESTION 7

- (Topic 1)

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

- A. True
- B. False

Answer: B

Explanation:

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

NEW QUESTION 8

- (Topic 1)

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

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

Answer: B

Explanation:

Snowflake's Time Travel feature is used for both querying historical data in tables and restoring and cloning historical data in databases, schemas, and tables³.

It allows users to access historical data within a defined period (1 day by default, up to 90 days for Snowflake Enterprise Edition) and is a key feature for data recovery and management. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 9

- (Topic 1)

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

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

Answer: BE

Explanation:

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

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Snowflake Documentation on Dynamic Data Masking1

NEW QUESTION 10

- (Topic 1)

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

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

Answer: BDE

Explanation:

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

NEW QUESTION 10

- (Topic 1)

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

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

Answer: CE

Explanation:

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

NEW QUESTION 14

- (Topic 1)

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

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

Answer: D

Explanation:

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

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

References:

- ? Snowflake Documentation on Roles
- ? Snowflake Documentation on Materialized Views

NEW QUESTION 18

- (Topic 1)

What happens when a cloned table is replicated to a secondary database? (Select TWO)

- A. A read-only copy of the cloned tables is stored.
- B. The replication will not be successful.
- C. The physical data is replicated
- D. Additional costs for storage are charged to a secondary account
- E. Metadata pointers to cloned tables are replicated

Answer: CE

Explanation:

When a cloned table is replicated to a secondary database in Snowflake, the following occurs:

? C. The physical data is replicated: The actual data of the cloned table is physically

replicated to the secondary database. This ensures that the secondary database has its own copy of the data, which can be used for read-only purposes or failover scenarios¹.

? E. Metadata pointers to cloned tables are replicated: Along with the physical data,

the metadata pointers that refer to the cloned tables are also replicated. This metadata includes information about the structure of the table and any associated properties².

It's important to note that while the physical data and metadata are replicated, the secondary database is typically read-only and cannot be used for write operations. Additionally, while there may be additional storage costs associated with the secondary account, this is not a direct result of the replication process but rather a consequence of storing additional data.

References:

? SnowPro Core Exam Prep — Answers to Snowflake's LEVEL UP: Backup and Recovery

? Snowflake SnowPro Core Certification Exam Questions Set 10

NEW QUESTION 19

- (Topic 1)

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

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

Answer: BD

Explanation:

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

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

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

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

References:

? Snowflake Documentation on Semi-Structured Data Types

? SnowPro® Core Certification Study Guide

NEW QUESTION 21

- (Topic 1)

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

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

Answer: AE

Explanation:

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

References:

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

? Snowflake's official documentation on billing and usage¹

NEW QUESTION 23

- (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 25

- (Topic 1)

Which services does the Snowflake Cloud Services layer manage? (Select TWO).

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

Answer: CE

Explanation:

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

References:

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

? Snowflake Documentation12 <https://docs.snowflake.com/en/user-guide/intro-key-concepts.html>

NEW QUESTION 29

- (Topic 1)

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

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

Answer: C

Explanation:

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

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

? Snowflake Documentation on Snowpipe1

NEW QUESTION 30

- (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 31

- (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 34

- (Topic 1)

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

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

Answer: B

Explanation:

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

NEW QUESTION 35

- (Topic 1)

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

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

Answer: AC

Explanation:

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

NEW QUESTION 37

- (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 38

- (Topic 1)

How often are encryption keys automatically rotated by Snowflake?

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

Answer: A

Explanation:

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

References:

- ? Understanding Encryption Key Management in Snowflake

NEW QUESTION 40

- (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 43

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

Answer: D

Explanation:

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

References:

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

NEW QUESTION 46

- (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 50

- (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 55

- (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 58

- (Topic 1)

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

- A. True
- B. False

Answer: A

Explanation:

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

NEW QUESTION 62

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 67

- (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 70

- (Topic 1)

What is the purpose of an External Function?

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

Answer: A

Explanation:

The purpose of an External Function in Snowflake is to call code that executes outside of the Snowflake environment. This allows Snowflake to interact with external services and leverage functionalities that are not natively available within Snowflake, such as calling APIs or running custom code hosted on cloud services. <https://docs.snowflake.com/en/sql-reference/external-functions.html>

NEW QUESTION 74

- (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 78

- (Topic 1)

In which use cases does Snowflake apply egress charges?

A. Data sharing within a specific region

B. Query result retrieval

C. Database replication

D. Loading data into Snowflake

Answer: C

Explanation:

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

References:

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

? Snowflake Documentation on Data Replication and Egress Charges¹

NEW QUESTION 83

- (Topic 1)

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

A. COPY INTO

B. GET

C. WRITE

D. EXTRACT INTO

Answer: A

Explanation:

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

References:

? Snowflake Documentation on Unloading Data

? Snowflake SnowPro Core: Copy Into Command to Unload Rows to Files in Named Stage

NEW QUESTION 85

- (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 87

- (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 92

- (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 93

- (Topic 1)

What transformations are supported in a CREATE PIPE ... AS COPY ... FROM (....) statement? (Select TWO.)

- A. Data can be filtered by an optional where clause
- B. Incoming data can be joined with other tables
- C. Columns can be reordered
- D. Columns can be omitted
- E. Row level access can be defined

Answer: AD

Explanation:

In a CREATE PIPE ... AS COPY ... FROM (....) statement, the supported transformations include filtering data using an optional WHERE clause and omitting columns. The WHERE clause allows for the specification of conditions to filter the data that is being loaded, ensuring only relevant data is inserted into the table. Omitting columns enables the exclusion of certain columns from the data load, which can be useful when the incoming data contains more columns than are needed for the target table.

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Simple Transformations During a Load¹

NEW QUESTION 98

- (Topic 1)

Which of the following are benefits of micro-partitioning? (Select TWO)

- A. Micro-partitions cannot overlap in their range of values
- B. Micro-partitions are immutable objects that support the use of Time Travel.
- C. Micro-partitions can reduce the amount of I/O from object storage to virtual warehouses
- D. Rows are automatically stored in sorted order within micro-partitions
- E. Micro-partitions can be defined on a schema-by-schema basis

Answer: BC

Explanation:

Micro-partitions in Snowflake are immutable objects, which means once they are written, they cannot be modified. This immutability supports the use of Time Travel, allowing users to access historical data within a defined period. Additionally, micro-partitions can significantly reduce the amount of I/O from object storage to virtual warehouses. This is because Snowflake's query optimizer can skip over micro-partitions that do not contain relevant data for a query, thus reducing the amount of data that needs to be scanned and transferred.

References: [COF-C02] SnowPro Core Certification Exam Study Guide <https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions.html>

NEW QUESTION 102

- (Topic 1)

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

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

Answer: C

Explanation:

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

NEW QUESTION 105

- (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 109

- (Topic 1)

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

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

Answer: CE

Explanation:

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

References:

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

? Snowflake Documentation on Snowpipe and Virtual Warehouses1

NEW QUESTION 113

- (Topic 1)

In the query profiler view for a query, which components represent areas that can be used to help optimize query performance? (Select TWO)

- A. Bytes scanned
- B. Bytes sent over the network
- C. Number of partitions scanned
- D. Percentage scanned from cache
- E. External bytes scanned

Answer: AC

Explanation:

In the query profiler view, the components that represent areas that can be used to help optimize query performance include ??Bytes scanned?? and ??Number of partitions scanned??. ??Bytes scanned?? indicates the total amount of data the query had to read and is a direct indicator of the query's efficiency. Reducing the bytes scanned can lead to lower data transfer costs and faster query execution. ??Number of partitions scanned?? reflects how well the data is clustered; fewer partitions scanned typically means better performance because the system can skip irrelevant data more effectively.

References:

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

? Snowflake Documentation on Query Profiling1

NEW QUESTION 118

- (Topic 1)

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

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

Answer: D

Explanation:

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

References:

- ? Snowflake Documentation on Clustering Keys & Clustered Tables1.
- ? Community discussions on how source data??s ordering affects a table with a cluster key

NEW QUESTION 119

- (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 121

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 125

- (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 130

- (Topic 1)

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

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

Answer: B

Explanation:

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

References:

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

? Snowflake Documentation on Clustering1

NEW QUESTION 133

- (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. DESCRIBE TABLE

Answer: A

Explanation:

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

References:

- ? Snowflake Documentation on Reader Accounts
- ? SnowPro® Core Certification Study Guide

NEW QUESTION 134

- (Topic 1)

What is a key feature of Snowflake architecture?

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

Answer: C

Explanation:

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

References:

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

NEW QUESTION 138

- (Topic 1)

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

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

Answer: AD

Explanation:

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

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

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

NEW QUESTION 141

- (Topic 2)

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

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

Answer: A

Explanation:

When a pipe is recreated using the CREATE OR REPLACE

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

NEW QUESTION 145

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 146

- (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 accounts.

NEW QUESTION 150

- (Topic 2)

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

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

Answer: AD

Explanation:

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

NEW QUESTION 152

- (Topic 2)

A user is preparing to load data from an external stage

Which practice will provide the MOST efficient loading performance?

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

Answer: A

Explanation:

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

NEW QUESTION 155

- (Topic 2)

What is the SNOWFLAKE.ACCOUNT_USAGE view that contains information about which objects were read by queries within the last 365 days (1 year)?

- A. VIEWS_HISTORY
- B. OBJECT_HISTORY
- C. ACCESS_HISTORY
- D. LOGIN_HISTORY

Answer: C

Explanation:

The ACCESS_HISTORY view in the SNOWFLAKE.ACCOUNT_USAGE schema contains information about the access history of Snowflake objects, such as tables and views, within the last 365 days.

NEW QUESTION 158

- (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 161

- (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 165

- (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 170

- (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 174

- (Topic 2)

Which of the following describes a Snowflake stored procedure?

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

Answer: D

Explanation:

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

NEW QUESTION 177

- (Topic 2)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 179

- (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 184

- (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 187

- (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 192

- (Topic 2)

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

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

Answer: BD

Explanation:

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

NEW QUESTION 196

- (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 198

- (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 200

- (Topic 2)

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

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

Answer: CE

Explanation:

Snowflake virtual warehouses support a configuration file and command line options in SnowSQL to specify a default warehouse, which is characteristic C. Additionally, the size of a virtual warehouse can be changed at any time, which is characteristic E. These features provide flexibility and ease of use in managing compute resources².
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

NEW QUESTION 204

- (Topic 2)

Which of the following is a data tokenization integration partner?

- A. Protegrity
- B. Tableau
- C. DBeaver
- D. SAP

Answer: A

Explanation:

Protegrity is listed as a data tokenization integration partner for Snowflake. This partnership allows Snowflake users to utilize Protegrity's tokenization solutions within the Snowflake environment³.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

NEW QUESTION 206

- (Topic 2)

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 210

- (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 211

- (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 212

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 215

- (Topic 2)

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

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

Answer: AD

Explanation:

Best practices for using the ACCOUNTADMIN role include ensuring that all users with this role use Multi-factor Authentication (MFA) for added security. Additionally, it is recommended to assign the ACCOUNTADMIN role to at least two users to avoid delays in case of password recovery issues, but to as few users as possible to maintain strict control over account-level operations⁴.

NEW QUESTION 218

- (Topic 2)

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

- A. Identifying queries that will likely run very slowly before executing them
- B. Locating queries that consume a high amount of credits
- C. Identifying logical issues with the queries
- D. Identifying inefficient micro-partition pruning

E. Data spilling to a local or remote disk

Answer: DE

Explanation:

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

NEW QUESTION 223

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 228

- (Topic 2)

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

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 230

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 235

- (Topic 2)

What is the MINIMUM edition of Snowflake that is required to use a SCIM security integration?

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

Answer: D

Explanation:

The minimum edition of Snowflake required to use a SCIM security integration is the Enterprise Edition. SCIM integrations are used for automated management of user identities and groups, and this feature is available starting from the Enterprise Edition of Snowflake. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 238

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 241

- (Topic 2)

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

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

Answer: AE

Explanation:

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

NEW QUESTION 246

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 248

- (Topic 2)

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

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

Answer: ABD

Explanation:

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

NEW QUESTION 249

- (Topic 2)

In an auto-scaling multi-cluster virtual warehouse with the setting SCALING_POLICY = ECONOMY enabled, when is another cluster started?

- A. When the system has enough load for 2 minutes
- B. When the system has enough load for 6 minutes
- C. When the system has enough load for 8 minutes
- D. When the system has enough load for 10 minutes

Answer: A

Explanation:

In an auto-scaling multi-cluster virtual warehouse with the SCALING_POLICY set to ECONOMY, another cluster is started when the system has enough load for 2 minutes (A). This policy is designed to optimize the balance between performance and cost, starting additional clusters only when the sustained load justifies it.

NEW QUESTION 251

- (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 255

- (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 258

- (Topic 2)

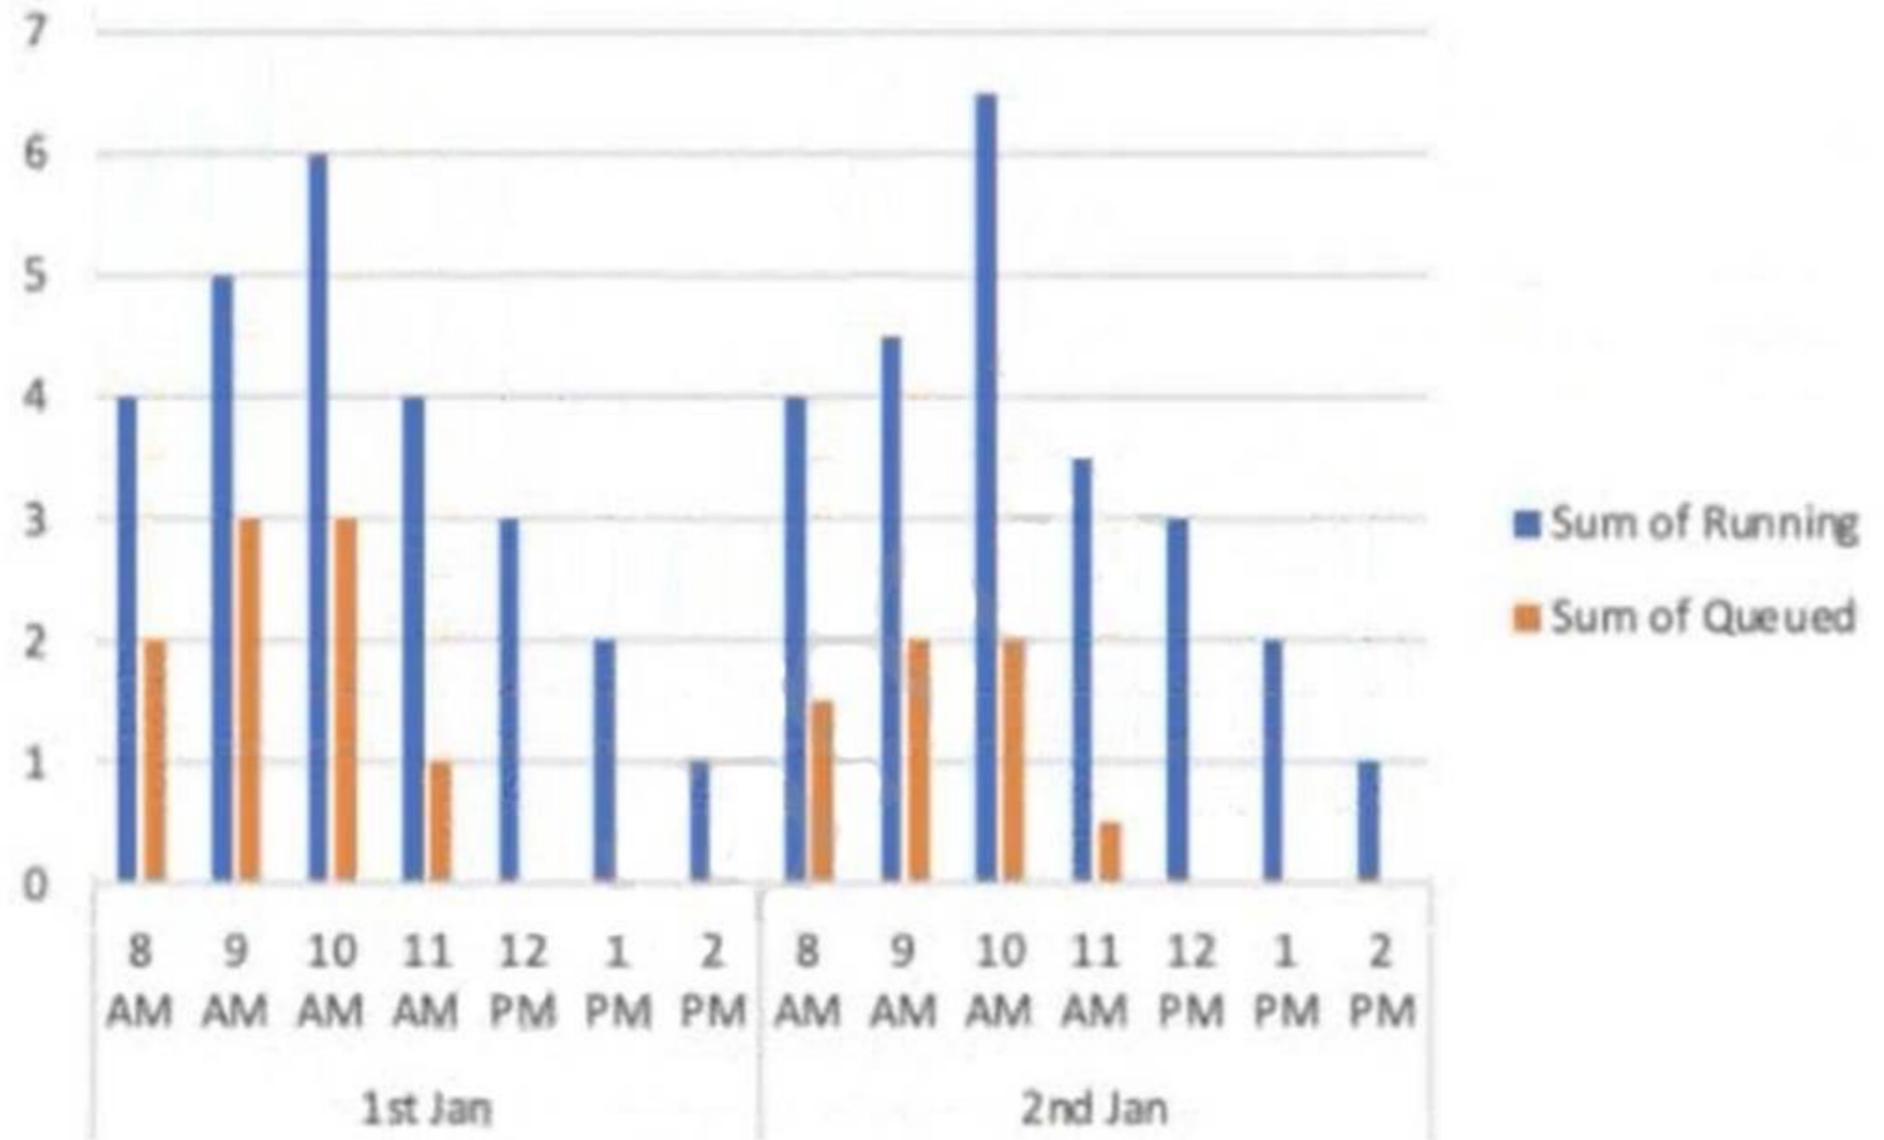
A virtual warehouse is created using the following command:

```
Create warehouse my_WH with warehouse_size = MEDIUM min_cluster_count = 1
```

```
max_cluster_count = 1
```

```
auto_suspend = 60 auto_resume = true;
```

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



What action should be taken to address this situation?

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

Answer: C

Explanation:

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

References = Snowflake Documentation on Multi-cluster Warehouses¹

NEW QUESTION 260

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 265

- (Topic 2)

Which minimum Snowflake edition allows for a dedicated metadata store?

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

Answer: B

Explanation:

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

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

NEW QUESTION 268

- (Topic 2)

A running virtual warehouse is suspended.

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

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

Answer: B

Explanation:

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

NEW QUESTION 271

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 275

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 277

- (Topic 2)

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

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

Answer: AD

Explanation:

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

NEW QUESTION 279

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 284

- (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 285

- (Topic 2)

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

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

Answer: AC

Explanation:

The SQL commands that consume a stream and advance the stream offset are those that result in changes to the data, such as UPDATE and INSERT operations. Specifically, `UPDATE TABLE FROM STREAM` and `INSERT INTO TABLE SELECT FROM STREAM` will consume the stream and move the offset forward, reflecting the changes made to the data. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 288

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 290

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

NEW QUESTION 292

- (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 294

- (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 298

- (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 303

- (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 304

- (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 307

- (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 309

- (Topic 3)

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

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

Answer: BC

Explanation:

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

NEW QUESTION 314

- (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 316

- (Topic 3)

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

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

Answer: BC

Explanation:

To optimize query performance in Snowflake, users can cluster a table, which organizes the data in a way that minimizes the amount of data scanned during queries. Additionally, enabling the search optimization service can improve the performance of selective point lookup queries on large tables³⁴.

NEW QUESTION 319

- (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 321

- (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 325

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 330

- (Topic 3)

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

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

Answer: AC

Explanation:

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

NEW QUESTION 332

- (Topic 3)

What happens when a database is cloned?

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

Answer: A

Explanation:

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

NEW QUESTION 334

- (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 335

- (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 layer3

NEW QUESTION 338

- (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 Duo1

NEW QUESTION 341

- (Topic 3)

Which of the following are characteristics of security in Snowflake?

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

Answer: C

Explanation:

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

NEW QUESTION 343

- (Topic 3)

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

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

Answer: AE

Explanation:

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

NEW QUESTION 347

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 348

- (Topic 3)

What is a characteristic of the Snowflake Query Profile?

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

Answer: B

Explanation:

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

NEW QUESTION 349

- (Topic 3)

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

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

Answer: BE

Explanation:

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

NEW QUESTION 352

- (Topic 3)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 357

- (Topic 3)

Which privilege must be granted to a share to allow secure views the ability to reference data in multiple databases?

- A. CREATE_SHARE on the account
- B. SHARE on databases and schemas
- C. SELECT on tables used by the secure view
- D. REFERENCE_USAGE on databases

Answer: D

Explanation:

To allow secure views the ability to reference data in multiple databases, the REFERENCE_USAGE privilege must be granted on each database that contains objects referenced by the secure view. This privilege is necessary before granting the SELECT privilege on a secure view to a share.

NEW QUESTION 360

- (Topic 3)

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

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

Answer: BDF

Explanation:

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

NEW QUESTION 361

- (Topic 3)

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

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

Answer: AE

Explanation:

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

NEW QUESTION 363

- (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 368

- (Topic 3)

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

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

Answer: ACF

Explanation:

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

NEW QUESTION 370

- (Topic 3)

What is the recommended compressed file size range for continuous data loads using Snowpipe?

- A. 8-16 MB
- B. 16-24 MB
- C. 10-99 MB
- D. 100-250 MB

Answer: D

Explanation:

For continuous data loads using Snowpipe, the recommended compressed file size range is between 100-250 MB. This size range is suggested to optimize the number of parallel operations for a load and to avoid size limitations, ensuring efficient and cost-effective data loading.

NEW QUESTION 374

- (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 intervention.

NEW QUESTION 375

- (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 377

- (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 378

- (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 381

- (Topic 3)

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

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

Answer: DE

Explanation:

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

NEW QUESTION 386

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 390

- (Topic 3)

What does Snowflake recommend regarding database object ownership? (Select TWO).

- A. Create objects with ACCOUNTADMIN and do not reassign ownership.
- B. Create objects with SYSADMIN.
- C. Create objects with SECURITYADMIN to ease granting of privileges later.
- D. Create objects with a custom role and grant this role to SYSADMIN.
- E. Use only MANAGED ACCESS SCHEMAS for objects owned by ACCOUNTADMIN.

Answer: BD

Explanation:

Snowflake recommends creating objects with a role that has the necessary privileges and is not overly permissive. SYSADMIN is typically used for managing system-level objects and operations. Creating objects with a custom role and granting this role to SYSADMIN allows for more granular control and adherence to the principle of least privilege. References: Based on best practices for database object ownership and role management.

NEW QUESTION 391

- (Topic 3)

How can a Snowflake user access a JSON object, given the following table? (Select TWO).

- A. src:salesperson.name
- B. src:sa1esPerso
- C. name
- D. src:salesperson.Name
- E. SRC:salesperson.name
- F. SRC:salesperson.Name

Answer: AC

Explanation:

To access a JSON object in Snowflake, dot notation is used where the path to the object is specified after the column name containing the JSON data. Both lowercase and uppercase can be used for attribute names, so both ??name?? and ??Name?? are valid. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 392

- (Topic 3)

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

Which command will show all privileges granted to that schema?

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

Answer: B

Explanation:

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

NEW QUESTION 393

- (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 395

- (Topic 3)

Which features make up Snowflake's column level security? (Select TWO).

- A. Continuous Data Protection (CDP)
- B. Dynamic Data Masking
- C. External Tokenization
- D. Key pair authentication
- E. Row access policies

Answer: BC

Explanation:

Snowflake's column level security features include Dynamic Data Masking and External Tokenization. Dynamic Data Masking uses masking policies to selectively mask data at query time, while External Tokenization allows for the tokenization of data before loading it into Snowflake and detokenizing it at query runtime⁵.

NEW QUESTION 399

- (Topic 3)

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

- A. Pre-signed URL

- B. Scoped URL
- C. Signed URL
- D. File URL

Answer: A

Explanation:

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

NEW QUESTION 400

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 403

- (Topic 3)

What does Snowflake's search optimization service support?

- A. External tables
- B. Materialized views
- C. Tables and views that are not protected by row access policies
- D. Casts on table columns (except for fixed-point numbers cast to strings)

Answer: C

Explanation:

Snowflake's search optimization service supports tables and views that are not protected by row access policies. It is designed to improve the performance of certain types of queries on tables, including selective point lookup queries and queries on fields in VARIANT, OBJECT, and ARRAY (semi-structured) columns

NEW QUESTION 405

- (Topic 3)

How long can a data consumer who has a pre-signed URL access data files using Snowflake?

- A. Indefinitely
- B. Until the result_cache expires
- C. Until the retention_time is met
- D. Until the expiration time is exceeded

Answer: D

Explanation:

A data consumer who has a pre-signed URL can access data files using Snowflake until the expiration time is exceeded. The expiration time is set when the pre-signed URL is generated and determines how long the URL remains valid

NEW QUESTION 409

- (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 410

- (Topic 3)

At what levels can a resource monitor be configured? (Select TWO).

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

Answer: AE

Explanation:

Resource monitors in Snowflake can be configured at the account and virtual warehouse levels. They are used to track credit usage and control costs associated with running virtual warehouses. When certain thresholds are reached, resource monitors can trigger actions such as sending alerts or suspending warehouses to prevent excessive credit consumption. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 413

- (Topic 3)

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

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

Answer: DE

Explanation:

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

NEW QUESTION 418

- (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 420

- (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 425

- (Topic 3)

Where is Snowflake metadata stored?

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

Answer: C

Explanation:

Snowflake's architecture is divided into three layers: database storage, query processing, and cloud services. The metadata, which includes information about the structure of the data, the SQL operations performed, and the service-level policies, is stored in the cloud services layer. This layer acts as the brain of the Snowflake environment, managing metadata, query optimization, and transaction coordination.

NEW QUESTION 427

- (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 430

- (Topic 3)

Which Snowflake URL type is used by directory tables?

- A. File
- B. Pre-signed
- C. Scoped
- D. Virtual-hosted style

Answer: C

Explanation:

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

NEW QUESTION 435

- (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 438

- (Topic 3)

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

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

Answer: BE

Explanation:

A Snowflake Data Exchange allows groups of accounts to share data privately among the accounts (B), and it supports bidirectional sharing, meaning an account can be both a provider and a consumer of data (E). This facilitates secure and governed data collaboration within a selected group³.

NEW QUESTION 440

- (Topic 3)

A tabular User-Defined Function (UDF) is defined by specifying a return clause that contains which keyword?

- A. ROW_NUMBER
- B. TABLE
- C. TABULAR
- D. VALUES

Answer: B

Explanation:

In Snowflake, a tabular User-Defined Function (UDF) is defined with a return clause that includes the keyword `TABLE`. This indicates that the UDF will return a set of rows, which can be used in the FROM clause of a query. References: Based on my internal knowledge as of 2021.

NEW QUESTION 442

- (Topic 3)

What type of columns does Snowflake recommend to be used as clustering keys? (Select TWO).

- A. A VARIANT column
- B. A column with very low cardinality
- C. A column with very high cardinality
- D. A column that is most actively used in selective filters
- E. A column that is most actively used in join predicates

Answer: CD

Explanation:

Snowflake recommends using columns with very high cardinality and those that are most actively used in selective filters as clustering keys. High cardinality columns have a wide range of unique values, which helps in evenly distributing the data across micro-partitions. Columns used in selective filters help in pruning the number of micro-partitions to scan, thus improving query performance. References: Based on general database optimization principles.

NEW QUESTION 447

- (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 tables.

NEW QUESTION 452

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 456

- (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 457

- (Topic 3)

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

Which condition will trigger a cluster to shut-down?

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

Answer: D

Explanation:

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

NEW QUESTION 461

- (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 465

.....

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](#)