

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)

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 7

- (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 8

- (Topic 1)

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

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

Answer: BD

Explanation:

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

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

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

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

? Snowflake Documentation on Virtual Warehouses

? SnowPro® Core Certification Study Guide

NEW QUESTION 9

- (Topic 1)

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

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

Answer: D

Explanation:

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

References:

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

NEW QUESTION 10

- (Topic 1)

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

- A. True
- B. False

Answer: A

Explanation:

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

Reference:<https://docs.snowflake.com/en/user-guide/warehouses-tasks.html#effects-of-resizing-a-suspended-warehouse>

NEW QUESTION 10

- (Topic 1)

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

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

Answer: AE

Explanation:

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

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Snowflake's official documentation on billing and usage1

NEW QUESTION 15

- (Topic 1)

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

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

Answer: C

Explanation:

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

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Snowflake Documentation on Snowpipe1

NEW QUESTION 19

- (Topic 1)

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

- A. Columns can be aggregated
- B. Columns can be joined with an existing table

- C. Columns can be reordered
- D. Columns can be omitted
- E. Data can be loaded without the need to spin up a virtual warehouse

Answer: CD

Explanation:

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

NEW QUESTION 24

- (Topic 1)

Which 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 25

- (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 capabilities.

NEW QUESTION 29

- (Topic 1)

User-level network policies can be created by which of the following roles? (Select TWO).

- A. ROLEADMIN
- B. ACCOUNTADMIN
- C. SYSADMIN
- D. SECURITYADMIN
- E. USERADMIN

Answer: BD

Explanation:

User-level network policies in Snowflake can be created by roles with the necessary privileges to manage security and account settings. The ACCOUNTADMIN role has the highest level of privileges across the account, including the ability to manage network policies. The SECURITYADMIN role is specifically responsible for managing security objects within Snowflake, which includes the creation and management of network policies.

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Snowflake Documentation on Network Policies
- ? Section 1.3 - SnowPro Core Certification Study Guide

NEW QUESTION 34

- (Topic 1)

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

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

Answer: B

Explanation:

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

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Snowflake Documentation on Account Usage Schema1

NEW QUESTION 36

- (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 37

- (Topic 1)

Which Snowflake technique can be used to improve the performance of a query?

- A. Clustering
- B. Indexing
- C. Fragmenting
- D. Using INDEX HINTS

Answer: A

Explanation:

Clustering is a technique used in Snowflake to improve the performance of queries. It involves organizing the data in a table into micro-partitions based on the values of one or more columns. This organization allows Snowflake to efficiently prune non-relevant micro-partitions during a query, which reduces the amount of data scanned and improves query performance.

References:

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

NEW QUESTION 39

- (Topic 1)

What is a best practice after creating a custom role?

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

Answer: B

Explanation:

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

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Section 1.3 - SnowPro Core Certification Study Guide1

NEW QUESTION 42

- (Topic 1)

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

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

Answer: AB

Explanation:

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

? A. VARIANT: This is a universal data type that can store values of any other type, including structured and semi-structured types. It is particularly useful for handling JSON, Avro, ORC, Parquet, and XML data formats1.

? B. ARRAY: An array is a list of elements that can be of any data type, including VARIANT, and is used to handle semi-structured data that is naturally represented as a list1.

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

References:

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

NEW QUESTION 43

- (Topic 1)

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

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

Answer: C

Explanation:

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

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Snowflake Documentation on Virtual Warehouse Performance¹

NEW QUESTION 45

- (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 48

- (Topic 1)

A sales table FCT_SALES has 100 million records. The following Query was executed
SELECT COUNT (1) FROM FCT SALES;
How did Snowflake fulfill this query?

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

Answer: D

Explanation:

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

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

NEW QUESTION 51

- (Topic 1)

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

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

Answer: D

Explanation:

In Snowflake, when the account is located on Microsoft Azure, a valid source for an external stage can be an Azure container or a folder path within an Azure

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

NEW QUESTION 52

- (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 56

- (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 services3. <https://docs.snowflake.com/en/sql-reference/external-functions.html>

NEW QUESTION 57

- (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 59

- (Topic 1)

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

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

Answer: DE

Explanation:

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

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

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

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

References:

- ? Snowflake Documentation on Clustering Keys & Clustered Tables
- ? Snowflake Documentation on SYSTEM\$CLUSTERING_INFORMATION
- ? Stack Overflow discussion on cluster key selection in Snowflake

NEW QUESTION 62

- (Topic 1)

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

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

Answer: A

Explanation:

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

References:

? Snowflake Documentation on Pipes

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

NEW QUESTION 64

- (Topic 1)

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

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

Answer: AE

Explanation:

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

? Privileges: The user must have the appropriate privileges on the objects associated with the query. This ensures that only authorized users can access cached data.

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

NEW QUESTION 68

- (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 73

- (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 78

- (Topic 1)

A marketing co-worker has requested the ability to change a warehouse size on their medium virtual warehouse called mktg WH. Which of the following statements will accommodate this request?

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

Answer: B

Explanation:

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

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Snowflake Documentation on Access Control Privileges1

NEW QUESTION 80

- (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 81

- (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 86

- (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 87

- (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 90

- (Topic 1)

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

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

Answer: D

Explanation:

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

NEW QUESTION 95

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 96

- (Topic 2)

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

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

Answer: AD

Explanation:

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

NEW QUESTION 97

- (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 structures1.

NEW QUESTION 102

- (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 103

- (Topic 2)

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

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

Answer: BC

Explanation:

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

NEW QUESTION 104

- (Topic 2)

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

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

Answer: BC

Explanation:

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

NEW QUESTION 109

- (Topic 2)

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

How can this requirement be met?

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

Answer: B

Explanation:

Dynamic data masking is a feature in Snowflake that allows administrators to define masking policies to protect sensitive data. It enables partial visibility of the data to certain roles, such as functional managers, while hiding it from others, like end users.

NEW QUESTION 112

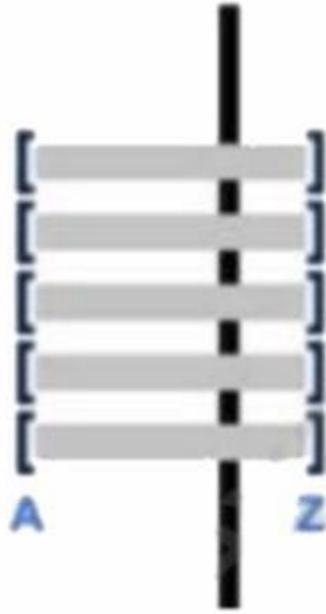
- (Topic 2)

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

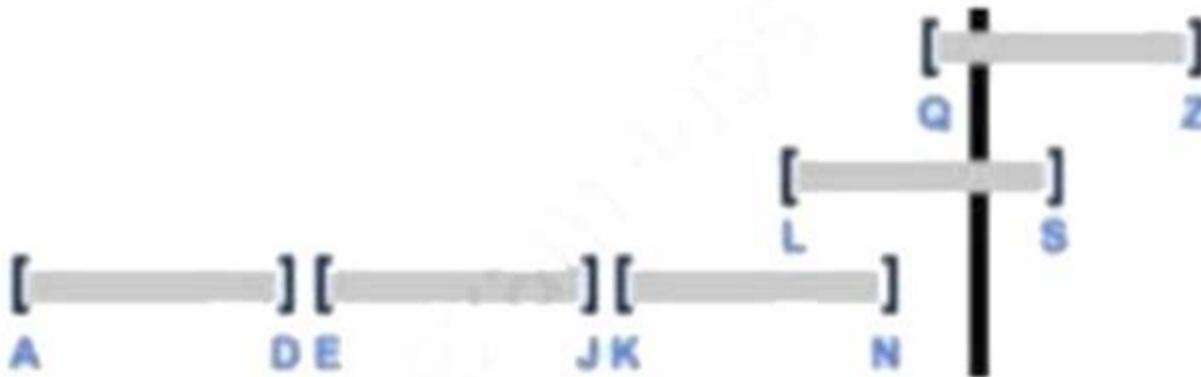
A.



B.



C.



D.



Answer: C

Explanation:

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

NEW QUESTION 114

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 118

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 123

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 124

- (Topic 2)

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

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

Answer: ACD

Explanation:

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

NEW QUESTION 127

- (Topic 2)

What is the purpose of multi-cluster virtual warehouses?

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

Answer: C

Explanation:

Multi-cluster virtual warehouses in Snowflake are designed to manage user and query concurrency needs. They allow for the allocation of additional clusters of compute resources, either statically or dynamically, to handle increased loads and reduce or eliminate the queuing of concurrent queries.
<https://docs.snowflake.com/en/user-guide/warehouses-multiclust.html#:~:text=Multi%2Dcluster%20warehouses%20enable%20you,during%20peak%20and%20off%20hours>.

NEW QUESTION 132

- (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 136

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 138

- (Topic 2)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 141

- (Topic 2)

How are serverless features billed?

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

Answer: B

Explanation:

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

NEW QUESTION 146

- (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 151

- (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 153

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 154

- (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 158

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 161

- (Topic 2)

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

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

Answer: BD

Explanation:

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

NEW QUESTION 164

- (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 167

- (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 171

- (Topic 2)

Which of the following describes the Snowflake Cloud Services layer?

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

Answer: A

Explanation:

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

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

NEW QUESTION 174

- (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 178

- (Topic 2)

How many days is load history for Snowpipe retained?

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

Answer: C

Explanation:

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

NEW QUESTION 182

- (Topic 2)

Which statement is true about running tasks in Snowflake?

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

Answer: B

Explanation:

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

NEW QUESTION 186

- (Topic 2)

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

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

Answer: BCF

Explanation:

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

NEW QUESTION 191

- (Topic 2)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 195

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 199

- (Topic 2)

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

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

Answer: B

Explanation:

Column-level security, which allows for the application of masking policies to columns in tables or views, is available starting from the Enterprise edition of Snowflake¹. References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation¹

NEW QUESTION 200

- (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 205

- (Topic 2)

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

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

Answer: ADE

Explanation:

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

NEW QUESTION 207

- (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 208

- (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 209

- (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 210

- (Topic 2)

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

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

Answer: BD

Explanation:

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

NEW QUESTION 213

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 215

- (Topic 2)

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

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

Answer: CE

Explanation:

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

NEW QUESTION 220

- (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 222

- (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 226

- (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 229

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 234

- (Topic 3)

Which parameter prevents streams on tables from becoming stale?

- A. MAXDATAEXTENSIONTIMEINDAYS
- B. MTN_DATA_RETENTION_TTIME_TN_DAYS
- C. LOCK_TIMEOUT
- D. STALE_AFTER

Answer: A

Explanation:

The parameter that prevents streams on tables from becoming stale is MAXDATAEXTENSIONTIMEINDAYS. This parameter specifies the maximum number of days for which Snowflake can extend the data retention period for the table to prevent streams on the table from becoming stale4.

NEW QUESTION 239

- (Topic 3)

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

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

D. #

Answer: C

Explanation:

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

NEW QUESTION 241

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 244

- (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 247

- (Topic 3)

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

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

Answer: C

Explanation:

The FLATTEN table function in Snowflake is used to convert semi- structured data, such as JSON or XML, into a relational format. It expands nested arrays or objects into multiple rows, making the data suitable for relational querying3.

NEW QUESTION 252

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 257

- (Topic 3)

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

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

D. 256

Answer: D

Explanation:

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

NEW QUESTION 260

- (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 263

- (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 264

- (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 268

- (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 271

- (Topic 3)

Which features could be used to improve the performance of queries that return a small subset of rows from a large table? (Select TWO).

- A. Search optimization service
- B. Automatic clustering

- C. Row access policies
- D. Multi-cluster virtual warehouses
- E. Secure views

Answer: AB

Explanation:

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

NEW QUESTION 276

- (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 278

- (Topic 3)

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

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

Answer: D

Explanation:

Snowflake retains information in the ACCESS HISTORY view for 365 days. This allows users to query the access history of Snowflake objects within the last year¹.

NEW QUESTION 283

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 284

- (Topic 3)

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 287

- (Topic 3)

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 289

- (Topic 3)

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

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

Answer: BDE

Explanation:

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

NEW QUESTION 294

- (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 295

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 298

- (Topic 3)

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

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

Answer: BC

Explanation:

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

NEW QUESTION 302

- (Topic 3)

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

How does Snowflake recommend maximizing the load performance?

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

Answer: B

Explanation:

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

[COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 303

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 308

- (Topic 3)

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

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

Answer: DE

Explanation:

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

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

NEW QUESTION 309

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 310

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 311

- (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 313

- (Topic 3)

When should a user consider disabling auto-suspend for a virtual warehouse? (Select TWO).

- A. When users will be using compute at different times throughout a 24/7 period
- B. When managing a steady workload
- C. When the compute must be available with no delay or lag time
- D. When the user does not want to have to manually turn on the warehouse each time it is needed
- E. When the warehouse is shared across different teams

Answer: BC

Explanation:

Disabling auto-suspend for a virtual warehouse is recommended when there is a steady workload, which ensures that compute resources are always available. Additionally, it is advisable to disable auto-suspend when immediate availability of compute resources is critical, eliminating any startup delay

NEW QUESTION 315

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 316

- (Topic 3)

Which statement describes pruning?

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

Answer: A

Explanation:

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

NEW QUESTION 321

- (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 322

- (Topic 3)

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

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

Answer: BC

Explanation:

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

NEW QUESTION 324

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 325

- (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 327

- (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 328

- (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 333

- (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 335

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 338

- (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 339

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 343

- (Topic 3)

What is the name of the SnowSQLfile 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 345

- (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 350

- (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 353

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 356

- (Topic 3)

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

What will occur?

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

Answer: A

Explanation:

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

NEW QUESTION 359

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 360

- (Topic 3)

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

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

Answer: BE

Explanation:

User-Defined Function (UDF) handlers must be written inline for Javascript and SQL. These languages allow the UDF logic to be included directly within the SQL statement that creates the UDF2.

NEW QUESTION 361

- (Topic 3)

Which REST API can be used with unstructured data?

- A. inscrtFilcs
- B. insertReport
- C. GET /api/tiles/
- D. loadHistoryScan

Answer: C

Explanation:

The REST API used with unstructured data in Snowflake is GET /api/files/, which retrieves (downloads) a data file from an internal or external stage4.

NEW QUESTION 366

- (Topic 4)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 369

- (Topic 4)

Which statistics can be used to identify queries that have inefficient pruning? (Select TWO).

- A. Bytes scanned
- B. Bytes written to result
- C. Partitions scanned
- D. Partitions total
- E. Percentage scanned from cache

Answer: CD

Explanation:

The statistics that can be used to identify queries with inefficient pruning are ??Partitions scanned?? and ??Partitions total??. These statistics indicate how much of the data was actually needed and scanned versus the total available, which can highlight inefficiencies in data pruning34.

NEW QUESTION 370

- (Topic 4)

How should clustering be used to optimize the performance of queries that run on a very large table?

- A. Manually re-cluster the table regularly.
- B. Choose one high cardinality column as the clustering key.
- C. Use the column that is most-frequently used in query select clauses as the clustering key.
- D. Assess the average table depth to identify how clustering is impacting the query.

Answer: B

Explanation:

For optimizing the performance of queries that run on a very large table, it is recommended to choose one high cardinality column as the clustering key. This helps to co-locate similar rows in the same micro-partitions, improving scan efficiency in queries by skipping data that does not match filtering predicates⁴.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 375

- (Topic 4)

What Snowflake feature provides a data hub for secure data collaboration, with a selected group of invited members?

- A. Data Replication
- B. Secure Data Sharing
- C. Data Exchange
- D. Snowflake Marketplace

Answer: C

Explanation:

Snowflake's Data Exchange feature provides a data hub for secure data collaboration. It allows providers to publish data that can be discovered and accessed by a selected group of invited members, facilitating secure and controlled data sharing within a collaborative environment³.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 377

- (Topic 4)

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

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

Answer: C

Explanation:

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

References:

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

NEW QUESTION 380

- (Topic 4)

What will prevent unauthorized access to a Snowflake account from an unknown source?

- A. Network policy
- B. End-to-end encryption
- C. Multi-Factor Authentication (MFA)
- D. Role-Based Access Control (RBAC)

Answer: A

Explanation:

A network policy in Snowflake is used to restrict access to the Snowflake account from unauthorized or unknown sources. It allows administrators to specify allowed IP address ranges, thus preventing access from any IP addresses not listed in the policy¹.

NEW QUESTION 381

- (Topic 4)

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

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

Answer: B

Explanation:

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

NEW QUESTION 385

- (Topic 4)

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

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

Answer: C

Explanation:

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

NEW QUESTION 388

- (Topic 4)

Which views are included in the DATA SHARING USAGE schema? (Select TWO).

- A. ACCESS_HISTORY
- B. DATA_TRANSFER_HISTORY
- C. WAREHOUSE_METERING_HISTORY
- D. MONETIZED_USAGE_DAILY
- E. LISTING_TELEMETRY_DAILY

Answer: DE

Explanation:

The DATA_SHARING_USAGE schema includes views that display information about listings published in the Snowflake Marketplace or a data exchange, which includes DATA_TRANSFER_HISTORY and LISTING_TELEMETRY_DAILY2.

NEW QUESTION 392

- (Topic 4)

A tag object has been assigned to a table (TABLE_A) in a schema within a Snowflake database.

Which CREATE object statement will automatically assign the TABLE_A tag to a target object?

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

Answer: C

Explanation:

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

NEW QUESTION 397

- (Topic 4)

Which Snowflake feature allows administrators to identify unused data that may be archived or deleted?

- A. Access history
- B. Data classification
- C. Dynamic Data Masking
- D. Object tagging

Answer: A

Explanation:

The Access History feature in Snowflake allows administrators to track data access patterns and identify unused data. This information can be used to make decisions about archiving or deleting data to optimize storage and reduce costs.

NEW QUESTION 399

- (Topic 4)

Which Snowflake feature provides increased login security for users connecting to Snowflake that is powered by Duo Security service?

- A. OAuth
- B. Network policies
- C. Single Sign-On (SSO)
- D. Multi-Factor Authentication (MFA)

Answer: D

Explanation:

Multi-Factor Authentication (MFA) provides increased login security for users connecting to Snowflake. Snowflake's MFA is powered by Duo Security service, which adds an additional layer of security during the login process.

NEW QUESTION 404

- (Topic 4)

How can performance be optimized for a query that returns a small amount of data from a very large base table?

- A. Use clustering keys
- B. Create materialized views
- C. Use the search optimization service
- D. Use the query acceleration service

Answer: C

Explanation:

The search optimization service in Snowflake is designed to improve the performance of selective point lookup queries on large tables, which is ideal for scenarios where a query returns a small amount of data from a very large base table¹. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 405

- (Topic 4)

Which type of loop requires a BREAK statement to stop executing?

- A. FOR
- B. LOOP
- C. REPEAT
- D. WHILE

Answer: B

Explanation:

The LOOP type of loop in Snowflake Scripting does not have a built-in termination condition and requires a BREAK statement to stop executing⁴.

NEW QUESTION 409

- (Topic 4)

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

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

Answer: AC

Explanation:

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

NEW QUESTION 410

- (Topic 4)

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

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

Answer: C

Explanation:

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

NEW QUESTION 411

- (Topic 4)

What does a masking policy consist of in Snowflake?

- A. A single data type, with one or more conditions, and one or more masking functions
- B. A single data type, with only one condition, and only one masking function
- C. Multiple data types, with only one condition, and one or more masking functions
- D. Multiple data types, with one or more conditions, and one or more masking functions

Answer: A

Explanation:

A masking policy in Snowflake consists of a single data type, with one or more conditions, and one or more masking functions. These components define how the data is masked based on the specified conditions³.

NEW QUESTION 414

- (Topic 4)

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

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

Answer: AC

Explanation:

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

NEW QUESTION 416

- (Topic 4)

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

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

Answer: A

Explanation:

The parameter DATA_RETENTION_TIME_IN_DAYS can be set at the account level to define the minimum number of days Snowflake retains historical data for Time Travel¹.

NEW QUESTION 417

- (Topic 4)

What does the LATERAL modifier for the FLATTEN function do?

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

Answer: C

Explanation:

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

NEW QUESTION 420

- (Topic 4)

Which privilege must be granted by one role to another role, and cannot be revoked?

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

Answer: C

Explanation:

The OWNERSHIP privilege is unique in that it must be granted by one role to another and cannot be revoked. This ensures that the transfer of ownership is deliberate and permanent, reflecting the importance of ownership in managing access and permissions.

NEW QUESTION 424

- (Topic 4)

When using the ALLOW_CLIENT_MFA_CACHING parameter, how long is a cached Multi-Factor Authentication (MFA) token valid for?

- A. 1 hour
- B. 2 hours
- C. 4 hours
- D. 8 hours

Answer: C

Explanation:

When using the ALLOW_CLIENT_MFA_CACHING parameter, a cached Multi-Factor Authentication (MFA) token is valid for up to 4 hours. This allows for continuous, secure connectivity without users needing to respond to an MFA prompt at the start of each connection attempt to Snowflake within this timeframe².

NEW QUESTION 425

- (Topic 4)

What function can be used with the recursive argument to return a list of distinct key names in all nested elements in an object?

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

Answer: A

Explanation:

The FLATTEN function can be used with the recursive argument to return a list of distinct key names in all nested elements within an object. This function is particularly useful for working with semi-structured data in Snowflake

NEW QUESTION 426

- (Topic 4)

Which command is used to unload data from a Snowflake database table into one or more files in a Snowflake stage?

- A. CREATE_STAGE
- B. COPY INTO <table>
- C. COPY INTO <location>
- D. CREATE PIPE

Answer: C

Explanation:

The COPY INTO <location> command is used to unload data from a Snowflake database table into one or more files in a Snowflake stage.

NEW QUESTION 429

- (Topic 4)

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

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

Answer: AE

Explanation:

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

NEW QUESTION 433

- (Topic 4)

A JSON file, that contains lots of dates and arrays, needs to be processed in Snowflake. The user wants to ensure optimal performance while querying the data. How can this be achieved?

- A. Flatten the data and store it in structured data types in a flattened table
- B. Query the table.
- C. Store the data in a table with a variant data type
- D. Query the table.
- E. Store the data in a table with a variant data type and include STRIP_NULL_VALUES while loading the table
- F. Query the table.
- G. Store the data in an external stage and create views on top of it
- H. Query the views.

Answer: B

Explanation:

Storing JSON data in a table with a VARIANT data type is optimal for querying because it allows Snowflake to leverage its semi-structured data capabilities. This approach enables efficient storage and querying without the need for flattening the data, which can be performance-intensive.

NEW QUESTION 434

- (Topic 4)

Which statistics are displayed in a Query Profile that indicate that intermediate results do not fit in memory? (Select TWO).

- A. Bytes scanned
- B. Partitions scanned
- C. Bytes spilled to local storage
- D. Bytes spilled to remote storage
- E. Percentage scanned from cache

Answer: CD

Explanation:

The Query Profile statistics that indicate intermediate results do not fit in memory are the bytes spilled to local storage and bytes spilled to remote storage.

NEW QUESTION 439

- (Topic 4)

What happens to the objects in a reader account when the DROP MANAGED ACCOUNT command is executed?

- A. The objects are dropped.
- B. The objects enter the Fail-safe period.
- C. The objects enter the Time Travel period.
- D. The objects are immediately moved to the provider account.

Answer: A

Explanation:

When the DROP MANAGED ACCOUNT command is executed in Snowflake, it removes the managed account, including all objects created within the account, and access to the account is immediately restricted².

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

NEW QUESTION 441

- (Topic 4)

What transformations are supported when loading data into a table using the COPY INTO <table> command? (Select TWO).

- A. Column reordering
- B. Column omission
- C. JOIN function
- D. FLATTEN function
- E. GROUP BY function

Answer: AB

Explanation:

The COPY INTO <table> command in Snowflake supports column reordering and column omission as part of its data transformation capabilities during the load process⁴⁵.

NEW QUESTION 444

- (Topic 4)

Which Snowflake role can manage any object grant globally, including modifying and revoking grants?

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

Answer: D

Explanation:

The SECURITYADMIN role in Snowflake can manage any object grant globally, including modifying and revoking grants. This role has the necessary privileges to oversee and control access to all securable objects within the Snowflake environment⁴.

NEW QUESTION 446

- (Topic 4)

What is a characteristic of materialized views in Snowflake?

- A. Materialized views do not allow joins.
- B. Clones of materialized views can be created directly by the user.
- C. Multiple tables can be joined in the underlying query of a materialized view.
- D. Aggregate functions can be used as window functions in materialized views.

Answer: C

Explanation:

One of the characteristics of materialized views in Snowflake is that they allow multiple tables to be joined in the underlying query. This enables the pre-computation of complex queries involving joins, which can significantly improve the performance of subsequent queries that access the materialized view⁴.

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

NEW QUESTION 448

- (Topic 5)

Which Snowflow object does not consume and storage costs?

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

Answer: C

Explanation:

Temporary tables in Snowflake do not consume storage costs. They are designed for transient data that is needed only for the duration of a session. Data stored in temporary tables is held in the virtual warehouse's cache and does not persist beyond the session's lifetime, thereby not incurring any storage charges.

References:

? Snowflake Documentation: Temporary Tables

NEW QUESTION 453

- (Topic 5)

The VALIDATE table function has which parameter as an input argument for a Snowflake user?

- A. Last_QUERY_ID
- B. CURRENT_STATEMENT
- C. UUID_STRING
- D. JOB_ID

Answer: C

Explanation:

The VALIDATE table function in Snowflake would typically use a unique identifier, such as a UUID_STRING, as an input argument. This function is designed to validate the data within a table against a set of constraints or conditions, often requiring a specific identifier to reference the particular data or job being validated.

References:

? There is no direct reference to a VALIDATE table function with these specific parameters in Snowflake documentation. It seems like a theoretical example for understanding function arguments. Snowflake documentation on UDFs and system functions can provide guidance on how to create and use custom functions for similar purposes.

NEW QUESTION 455

- (Topic 5)

Which URL provides access to files in Snowflake without authorization?

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

Answer: C

Explanation:

A Pre-signed URL provides access to files stored in Snowflake without requiring authorization at the time of access. This feature allows users to generate a URL with a limited validity period that grants temporary access to a file in a secure manner. It's particularly useful for sharing data with external parties or applications without the need for them to authenticate directly with Snowflake.

References:

? Snowflake Documentation: Using Pre-signed URLs

NEW QUESTION 456

- (Topic 5)

A Snowflake user wants to temporarily bypass a network policy by configuring the user object property MINS_TO_BYPASS_NETWORK_POLICY. What should they do?

- A. Use the SECURITYADMIN role.
- B. Use the SYSADMIN role.
- C. Use the USERADMIN role.
- D. Contact Snowflake Support.

Answer: C

Explanation:

To temporarily bypass a network policy by configuring the user object property MINS_TO_BYPASS_NETWORK_POLICY, the USERADMIN role should be used. This role has the necessary privileges to modify user properties, including setting a temporary bypass for network policies, which can be crucial for enabling access under specific circumstances without permanently altering the network security configuration. References:

? Snowflake Documentation: User Management

NEW QUESTION 457

- (Topic 5)

Which function will provide the proxy information needed to protect Snowsight?

- A. SYSTEMADMIN_TAG
- B. SYSTEM\$GET_PRIVATELINK
- C. SYSTEMSALLONTLIST
- D. SYSTEMAUTHORIZE

Answer: B

Explanation:

The SYSTEM\$GET_PRIVATELINK function in Snowflake provides proxy information necessary for configuring PrivateLink connections, which can protect Snowsight as well as other Snowflake services. PrivateLink enhances security by allowing Snowflake to be accessed via a private connection within a cloud provider's network, reducing exposure to the public internet.

References:

? Snowflake Documentation: PrivateLink Setup

NEW QUESTION 462

- (Topic 5)

What are the benefits of the replication feature in Snowflake? (Select TWO).

- A. Disaster recovery
- B. Time Travel
- C. Fail-safe
- D. Database failover and fallback
- E. Data security

Answer: AD

Explanation:

The replication feature in Snowflake provides several benefits, with disaster recovery and database failover and fallback being two of the primary advantages. Replication allows for the continuous copying of data from one Snowflake account to another, ensuring that a secondary copy of the data is available in case of outages or disasters. This capability supports disaster recovery strategies by allowing operations to quickly switch to the replicated data in a different account or region. Additionally, it facilitates database failover and fallback procedures, ensuring business continuity and minimizing downtime.

References:

? Snowflake Documentation: Data Replication

NEW QUESTION 465

- (Topic 5)

What is used to denote a pre-computed data set derived from a SELECT query specification and stored for later use?

- A. View
- B. Secure view
- C. Materialized view
- D. External table

Answer: C

Explanation:

A materialized view in Snowflake denotes a pre-computed data set derived from a SELECT query specification and stored for later use. Unlike standard views, which dynamically compute the data each time the view is accessed, materialized views store the result of the query at the time it is executed, thereby speeding up access to the data, especially for expensive aggregations on large datasets.

References:

? Snowflake Documentation: Materialized Views

NEW QUESTION 469

- (Topic 5)

How does a Snowflake stored procedure compare to a User-Defined Function (UDF)?

- A. A single executable statement can call only two stored procedure
- B. In contrast, a single SQL statement can call multiple UDFs.
- C. A single executable statement can call only one stored procedur
- D. In contrast, a single SQL statement can call multiple UDFs.
- E. A single executable statement can call multiple stored procedure
- F. In contrast, multiple SQL statements can call the same UDFs.
- G. Multiple executable statements can call more than one stored procedur
- H. In contrast, a single SQL statement can call multiple UDFs.

Answer: B

Explanation:

In Snowflake, stored procedures and User-Defined Functions (UDFs) have different invocation patterns within SQL:

? Option B is correct: A single executable statement can call only one stored procedure due to the procedural and potentially transactional nature of stored procedures. In contrast, a single SQL statement can call multiple UDFs because UDFs are designed to operate more like functions in traditional programming, where they return a value and can be embedded within SQL queries. References: Snowflake documentation comparing the operational differences between stored procedures and UDFs.

NEW QUESTION 473

.....

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](#)