

Snowflake

Exam Questions COF-C02

SnowPro Core Certification Exam (COF-C02)



NEW QUESTION 1

- (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 2

- (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 3

- (Topic 1)

What features does Snowflake Time Travel enable?

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

Answer: BC

Explanation:

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

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

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

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

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

While Time Travel does allow querying of past data, it is limited to the retention period set for the Snowflake account, which is typically 1 day for standard accounts and can be extended up to 90 days for enterprise accounts. It does not enable querying or restoring objects created or deleted beyond the retention period, nor does it provide analysis over all periods of time.

References:

? Snowflake Documentation on Time Travel

? SnowPro® Core Certification Study Guide

NEW QUESTION 4

- (Topic 1)

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 5

- (Topic 1)

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

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

Answer: CE

Explanation:

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

NEW QUESTION 6

- (Topic 1)

A company strongly encourages all Snowflake users to self-enroll in Snowflake's default Multi-Factor Authentication (MFA) service to provide increased login security for users connecting to Snowflake.

Which application will the Snowflake users need to install on their devices in order to connect with MFA?

- A. Okta Verify
- B. Duo Mobile
- C. Microsoft Authenticator
- D. Google Authenticator

Answer: B

Explanation:

Snowflake's default Multi-Factor Authentication (MFA) service is powered by Duo Security. Users are required to install the Duo Mobile application on their devices to

use MFA for increased login security when connecting to Snowflake. This service is managed entirely by Snowflake, and users do not need to sign up separately with Duo¹.

NEW QUESTION 7

- (Topic 1)

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

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

Answer: AE

Explanation:

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

References:

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

? Snowflake's official documentation on billing and usage¹

NEW QUESTION 8

- (Topic 1)

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

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

Answer: CE

Explanation:

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

References:

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

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

NEW QUESTION 9

- (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 10

- (Topic 1)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 10

- (Topic 1)

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

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

Answer: AC

Explanation:

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

NEW QUESTION 14

- (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 Policies¹

? Section 1.3 - SnowPro Core Certification Study Guide

NEW QUESTION 18

- (Topic 1)

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

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

Answer: DE

Explanation:

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

References:

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

? Snowflake Documentation on Secure Data Sharing¹

NEW QUESTION 21

- (Topic 1)

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

- A. Via the Snowflake Web Interface (UI)
- B. Via thedata_share=trueparameter
- C. Via SQL commands
- D. Via Virtual Warehouses

Answer: AC

Explanation:

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

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

? Via SQL commands: Snowflake also allows the creation and management of shares using SQL commands. This method is more suited for users who prefer scripting or need to automate the process.

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

NEW QUESTION 24

- (Topic 1)

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

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

Answer: C

Explanation:

Transient tables in Snowflake have a default Time Travel retention period of 1 day, which allows users to access historical data within the last 24 hours. However, transient tables do not have a Fail-safe period. Fail-safe is an additional layer of data protection that retains data beyond the Time Travel period for recovery purposes in case of extreme data loss. Since transient tables are designed for temporary or intermediate workloads with no requirement for long-term durability, they do not include a Fail-safe period by default¹.

References:

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

NEW QUESTION 27

- (Topic 1)

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

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

Answer: D

Explanation:

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

References:

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

? Snowflake Documentation on Snowflake Architecture¹

NEW QUESTION 32

- (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 35

- (Topic 1)

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

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

Answer: AB

Explanation:

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

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

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

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

References:

? Snowflake Documentation on Multi-Cluster Warehousing

? SnowPro Core Certification best practices

NEW QUESTION 36

- (Topic 1)

In which use cases does Snowflake apply egress charges?

- A. Data sharing within a specific region
- B. Query result retrieval
- C. Database replication
- D. Loading data into Snowflake

Answer: C

Explanation:

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

References:

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

? Snowflake Documentation on Data Replication and Egress Charges1

NEW QUESTION 37

- (Topic 1)

During periods of warehouse contention which parameter controls the maximum length of time a warehouse will hold a query for processing?

- A. STATEMENT_TIMEOUT IN SECONDS
- B. STATEMENT_QUEUED_TIMEOUT_IN_SECONDS
- C. MAX_CONCURRENCY LEVEL
- D. QUERY_TIMEOUT_IN_SECONDS

Answer: B

Explanation:

The parameter STATEMENT_QUEUED_TIMEOUT_IN_SECONDS sets the limit for a query to wait in the queue in order to get its chance of running on the warehouse. The query will quit after reaching this limit. By default, the value of this parameter is 0 which mean the queries will wait indefinitely in the waiting queue [https://community.snowflake.com/s/article/Warehouse-Concurrency-and-Statement- Timeout-Parameters#:~:text=The%20parameter%20STATEMENT_QUEUED_TIMEOUT_IN_SECO NDS%20sets%20the,indefinitely%20in%20the%20waiting%20queue.](https://community.snowflake.com/s/article/Warehouse-Concurrency-and-Statement-Timeout-Parameters#:~:text=The%20parameter%20STATEMENT_QUEUED_TIMEOUT_IN_SECONDS%20sets%20the,indefinitely%20in%20the%20waiting%20queue.)

NEW QUESTION 38

- (Topic 1)

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

- A. True
- B. False

Answer: B

Explanation:

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

resources in an active warehouse. References: Preparing your data files | Snowflake Documentation

NEW QUESTION 40

- (Topic 1)

Which statement about billing applies to Snowflake credits?

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

Answer: D

Explanation:

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

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

NEW QUESTION 45

- (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 46

- (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 48

- (Topic 2)

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

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

Answer: A

Explanation:

When a pipe is recreated using the CREATE OR REPLACE

PIPE command, the load history of the pipe is reset. This means that Snowpipe will consider all files in the stage as new and will attempt to load them, even if they were loaded previously by the old pipe².

NEW QUESTION 51

- (Topic 2)

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

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

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

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

Answer: B

Explanation:

In Snowflake, to extract a specific element from a JSON stored in a VARIANT column, the correct syntax is to use the dot notation. Therefore, the query select src.dealership from car_sales; will return the dealership information contained within each JSON object in the src column. References: For a detailed explanation, please refer to the Snowflake documentation on querying semi-structured data.

NEW QUESTION 52

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 53

- (Topic 2)

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

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

Answer: AD

Explanation:

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

NEW QUESTION 56

- (Topic 2)

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

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

Answer: C

Explanation:

The ACCESS_HISTORY view in the SNOWFLAKE.ACCOUNT_USAGE schema contains information about the access history of Snowflake objects, such as tables and views, within the last 365 days¹.

NEW QUESTION 57

- (Topic 2)

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

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

Answer: AC

Explanation:

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

NEW QUESTION 61

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 62

- (Topic 2)

When loading data into Snowflake via Snowpipe what is the compressed file size recommendation?

- A. 10-50 MB
- B. 100-250 MB
- C. 300-500 MB
- D. 1000-1500 MB

Answer: B

Explanation:

For loading data into Snowflake via Snowpipe, the recommended compressed file size is between 100-250 MB. This size range is optimal for balancing the performance of parallel processing and minimizing the overhead associated with handling many small files².

NEW QUESTION 64

- (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 67

- (Topic 2)

What affects whether the query results cache can be used?

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

Answer: C

Explanation:

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

NEW QUESTION 71

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 75

- (Topic 2)

Which command sets the Virtual Warehouse for a session?

- A. COPY WAREHOUSE FROM <<config file>>;
- B. SET WAREHOUSE = <<warehouse name>>;
- C. USE WAREHOUSE <<warehouse name>>;
- D. USE VIRTUAL_WAREHOUSE <<warehouse name>>;

Answer: C

Explanation:

The command USE WAREHOUSE <<warehouse name>>; is used to set the virtual warehouse for the current session in Snowflake. This command specifies which virtual warehouse to use for executing queries in that session¹.

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

NEW QUESTION 76

- (Topic 2)

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 81

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 85

- (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 87

- (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 92

- (Topic 2)

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

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

Answer: CE

Explanation:

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

NEW QUESTION 95

- (Topic 2)

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

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

Answer: B

Explanation:

Operations that do not require compute resources are typically those that can leverage previously cached results. However, if no queries have been executed previously, all the given operations would require compute to execute. It's important to note that certain operations like DDL statements and queries that hit the result cache do not consume compute credits².

NEW QUESTION 100

- (Topic 2)

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

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

Answer: AD

Explanation:

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

NEW QUESTION 104

- (Topic 2)

What is an advantage of using an explain plan instead of the query profiler to evaluate the performance of a query?

- A. The explain plan output is available graphically.
- B. An explain plan can be used to conduct performance analysis without executing a query.
- C. An explain plan will handle queries with temporary tables and the query profiler will not.
- D. An explain plan's output will display automatic data skew optimization information.

Answer: B

Explanation:

An explain plan is beneficial because it allows for the evaluation of how a query will be processed without the need to actually execute the query. This can help in understanding the query's performance implications and potential bottlenecks without consuming resources that would be used if the query were run

NEW QUESTION 106

- (Topic 2)

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

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 109

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 113

- (Topic 2)

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

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

Answer: B

Explanation:

The GET command is used to download files from a Snowflake stage to a local folder on a client's machine.

Reference: <https://docs.snowflake.com/en/sql-reference/sql/get.html>

NEW QUESTION 116

- (Topic 2)

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

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

Answer: ABD

Explanation:

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

NEW QUESTION 121

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 122

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 125

- (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 126

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 129

- (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 133

- (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 137

- (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 139

- (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 144

- (Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 148

- (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 153

- (Topic 3)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 154

- (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 155

- (Topic 3)

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

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

Answer: D

Explanation:

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

which allows for flexibility in the data structure

NEW QUESTION 156

- (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 160

- (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 161

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 164

- (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 168

- (Topic 3)

Which Snowflake edition enables data sharing only through Snowflake Support?

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

Answer: A

Explanation:

The Snowflake edition that enables data sharing only through Snowflake Support is the Virtual Private Snowflake (VPS). By default, VPS does not permit data

sharing outside of the VPS environment, but it can be enabled through Snowflake Support4.

NEW QUESTION 172

- (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 176

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 178

- (Topic 3)

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

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

Answer: AE

Explanation:

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

NEW QUESTION 181

- (Topic 3)

What is a characteristic of the Snowflake Query Profile?

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

Answer: B

Explanation:

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

NEW QUESTION 186

- (Topic 3)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 188

- (Topic 3)

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

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

Answer: D

Explanation:

To allow secure views the ability to reference data in multiple databases, the REFERENCE_USAGE privilege must be granted on each database that contains objects referenced by the secure view². This privilege is necessary before granting the SELECT privilege on a secure view to a share.

NEW QUESTION 189

- (Topic 3)

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

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

Answer: BDF

Explanation:

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

NEW QUESTION 193

- (Topic 3)

What computer language can be selected when creating User-Defined Functions (UDFs) using the Snowpark API?

- A. Swift
- B. JavaScript
- C. Python
- D. SQL

Answer: C

Explanation:

The Snowpark API allows developers to create User-Defined Functions (UDFs) in various languages, including Python, which is known for its ease of use and wide adoption in data-related tasks. References: Based on general programming and cloud data service knowledge as of 2021.

NEW QUESTION 194

- (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 199

- (Topic 3)

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

- A. Include the SQL statements in the body of the task CREATE TASK mytask .. AS INSERT INTO target1 SELECT .. FROM stream_s1 WHERE .. INSERT INTO target2 SELECT .. FROM stream_s1 WHERE ..
- B. A stored procedure can have only one DML statement per stored procedure invocation and therefore the user should sequence stored procedure calls in the taskdefinition CREATE TASK mytask AS call stored_proc1(); call stored_proc2();

- C. Use a stored procedure executing multiple SQL statements and invoke the stored procedure from the task
- D. CREATE TASK mytask AS call stored_proc_multiple_statements_inside();
- E. Create a task for each SQL statement (e.
- F. resulting in task1, task2, etc.) and string the series of SQL statements by having a control task calling task1, task2, etc.
- 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 204

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 207

- (Topic 3)

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

Which command will show all privileges granted to that schema?

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

Answer: B

Explanation:

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

NEW QUESTION 208

- (Topic 3)

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

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

Answer: B

Explanation:

In Snowflake, the OPERATE privilege is required for a role to suspend or resume a task. This privilege allows the role to perform operational tasks such as starting and stopping tasks, which includes suspending and resuming them⁶

NEW QUESTION 211

- (Topic 3)

What does Snowflake's search optimization service support?

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

Answer: C

Explanation:

Snowflake's search optimization service supports tables and views that are not protected by row access policies. It is designed to improve the performance of certain types of queries on tables, including selective point lookup queries and queries on fields in VARIANT, OBJECT, and ARRAY (semi-structured) columns¹.

NEW QUESTION 213

- (Topic 3)

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

- * 1. Created a user called "CONSUMER"
 - * 2. Created a database to hold the share and an extra-small warehouse to query the data
 - * 3. Granted the role PUBLIC the following privileges: Usage on the warehouse, database, and schema, and SELECT on all the objects in the share
- Based on this configuration what is true of the reader account?

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

Answer: B

Explanation:

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

Very Comprehensive Explanation

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

NEW QUESTION 218

- (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 221

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 226

- (Topic 3)

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

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

Answer: DE

Explanation:

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

NEW QUESTION 231

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 233

- (Topic 3)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 235

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 238

- (Topic 3)

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

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

Answer: BD

Explanation:

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

NEW QUESTION 240

- (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 244

- (Topic 3)

What is cached during a query on a virtual warehouse?

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

Answer: C

Explanation:

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

NEW QUESTION 249

- (Topic 3)

How does Snowflake recommend handling the bulk loading of data batches from files already available in cloud storage?

- A. Use Snowpipe.
- B. Use the INSERT command.
- C. Use an external table.
- D. Use the COPY command.

Answer: D

Explanation:

Snowflake recommends using the COPY command for bulk loading data batches from files already available in cloud storage. This command allows for efficient and large-scale data loading operations from files staged in cloud storage into Snowflake tables3.

NEW QUESTION 254

- (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 257

- (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 262

- (Topic 4)

Which Snowflake feature allows a user to track sensitive data for compliance, discovery, protection, and resource usage?

- A. Tags
- B. Comments
- C. Internal tokenization
- D. Row access policies

Answer: A

Explanation:

Tags in Snowflake allow users to track sensitive data for compliance, discovery, protection, and resource usage. They enable the categorization and tracking of data, supporting compliance with privacy regulations678. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 263

- (Topic 4)

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

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

Answer: C

Explanation:

To determine which user has accessed a database object containing sensitive information, a Snowflake administrator can query the ACCESS_HISTORY view in the ACCOUNT_USAGE schema, which provides information about access to database objects³.

NEW QUESTION 265

- (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 267

- (Topic 4)

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

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

Answer: AE

Explanation:

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

NEW QUESTION 271

- (Topic 4)

Which function unloads data from a relational table to JSON?

- A. TO_OBJECT
- B. TO_JSON
- C. TO_VARIANT
- D. OBJECT CONSTRUCT

Answer: B

Explanation:

The TO_JSON function is used to convert a VARIANT value into a string containing the JSON representation of the value. This function is suitable for unloading data from a relational table to JSON format. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 274

- (Topic 4)

What metadata does Snowflake store for rows in micro-partitions? (Select TWO).

- A. Range of values
- B. Distinct values
- C. Index values
- D. Sorted values
- E. Null values

Answer: AB

Explanation:

Snowflake stores metadata for rows in micro-partitions, including the range of values for each column and the number of distinct values¹.

NEW QUESTION 275

- (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 277

- (Topic 4)

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

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

Answer: D

Explanation:

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

NEW QUESTION 282

- (Topic 4)

How is unstructured data retrieved from data storage?

- A. SQL functions like the GET command can be used to copy the unstructured data to a location on the client.
- B. SQL functions can be used to create different types of URLs pointing to the unstructured data
- C. These URLs can be used to download the data to a client.
- D. SQL functions can be used to retrieve the data from the query results cache
- E. When the query results are output to a client, the unstructured data will be output to the client as files.
- F. SQL functions can call on different web extensions designed to display different types of files as a web page
- G. The web extensions will allow the files to be downloaded to the client.

Answer: B

Explanation:

Unstructured data stored in Snowflake can be retrieved by using SQL functions to generate URLs that point to the data. These URLs can then be used to download the data directly to a client

NEW QUESTION 284

- (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 285

- (Topic 4)

How can a Snowflake user traverse semi-structured data?

- A. Insert a colon (:) between the VARIANT column name and any first-level element.
- B. Insert a colon (:) between the VARIANT column name and any second-level element
- C. Insert a double colon (: :) between the VARIANT column name and any first-level element.
- D. Insert a double colon (: :) between the VARIANT column name and any second-level element.

Answer: A

Explanation:

To traverse semi-structured data in Snowflake, a user can insert a colon (:) between the VARIANT column name and any first-level element. This path syntax is used to retrieve elements in a VARIANT column⁴.

NEW QUESTION 286

- (Topic 4)

Which object can be used with Secure Data Sharing?

- A. View

- B. Materialized view
- C. External table
- D. User-Defined Function (UDF)

Answer: A

Explanation:

Views can be used with Secure Data Sharing in Snowflake. Materialized views, external tables, and UDFs are not typically shared directly for security and performance reasons².

NEW QUESTION 288

- (Topic 4)

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

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

Answer: C

Explanation:

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

NEW QUESTION 291

- (Topic 4)

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

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

Answer: D

Explanation:

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

NEW QUESTION 292

- (Topic 4)

Which ACCOUNT_USAGE schema database role provides visibility into policy-related information?

- A. USAGE_VIEWER
- B. GOVERNANCE_VIEWER
- C. OBJECT_VIEWER
- D. SECURITY_VIEWER

Answer: B

Explanation:

The GOVERNANCE_VIEWER role in the ACCOUNT_USAGE schema provides visibility into policy-related information within Snowflake. This role is specifically designed to access views that display object metadata and usage metrics related to governance¹².

NEW QUESTION 295

- (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 299

- (Topic 4)

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

- A. Temporary tables
- B. Permanent tables

- C. Transient tables
- D. External tables
- E. User-Defined Table Functions (UDTFs)

Answer: BD

Explanation:

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

NEW QUESTION 303

- (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 306

- (Topic 4)

Who can grant object privileges in a regular schema?

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

Answer: A

Explanation:

In a regular schema within Snowflake, the object owner has the privilege to grant object privileges. The object owner is typically the role that created the object or to whom the ownership of the object has been transferred⁷⁸.

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

NEW QUESTION 308

- (Topic 4)

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

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

Answer: C

Explanation:

To start configuring Snowflake for Single Sign-On (SSO), the CREATE SECURITY INTEGRATION command is used. This command sets up a security integration object in Snowflake, which is necessary for enabling SSO with external identity providers using SAML 2.0¹.

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

NEW QUESTION 310

- (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 311

- (Topic 4)

While working with unstructured data, which file function generates a Snowflake-hosted file URL to a staged file using the stage name and relative file path as inputs?

- A. GET_PREIGNED_URL
- B. GET_ABSOLUTE_PATH
- C. BUILD_STAGE_FILE_URL
- D. BUILD SCOPED FILE URL

Answer: C

Explanation:

The BUILD_STAGE_FILE_URL function generates a Snowflake-hosted file URL to a staged file using the stage name and relative file path as inputs².

NEW QUESTION 312

- (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 stage¹.

NEW QUESTION 314

- (Topic 4)

What is the purpose of the STRIP NULL_VALUES file format option when loading semi- structured data files into Snowflake?

- A. It removes null values from all columns in the data.
- B. It converts null values to empty strings during loading.
- C. It skips rows with null values during the loading process.
- D. It removes object or array elements containing null values.

Answer: D

Explanation:

The STRIP NULL_VALUES file format option, when set to TRUE, removes object or array elements that contain null values during the loading process of semi-structured data files into Snowflake. This ensures that the data loaded into Snowflake tables does not contain these null elements, which can be useful when the ??null?? values in files indicate missing values and have no other special meaning².

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

NEW QUESTION 318

- (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 320

- (Topic 4)

What does SnowCD help Snowflake users to do?

- A. Copy data into files.
- B. Manage different databases and schemas.
- C. Troubleshoot network connections to Snowflake.
- D. Write SELECT queries to retrieve data from external tables.

Answer: C

Explanation:

SnowCD is a connectivity diagnostic tool that helps users troubleshoot network connections to Snowflake. It performs a series of checks to evaluate the network connection and provides suggestions for resolving any issues⁴.

NEW QUESTION 325

- (Topic 4)

Which Snowflake function is maintained separately from the data and helps to support features such as Time Travel, Secure Data Sharing, and pruning?

- A. Column compression

- B. Data clustering
- C. Micro-partitioning
- D. Metadata management

Answer: C

Explanation:

Micro-partitioning is a Snowflake function that is maintained separately from the data and supports features such as Time Travel, Secure Data Sharing, and pruning. It allows Snowflake to efficiently manage and query large datasets by organizing them into micro-partitions¹.

NEW QUESTION 329

- (Topic 4)

What does a Notify & Suspend action for a resource monitor do?

- A. Send an alert notification to all account users who have notifications enabled.
- B. Send an alert notification to all virtual warehouse users when thresholds over 100% have been met.
- C. Send a notification to all account administrators who have notifications enabled, and suspend all assigned warehouses after all statements being executed by the warehouses have completed.
- D. Send a notification to all account administrators who have notifications enabled, and suspend all assigned warehouses immediately, canceling any statements being executed by the warehouses.

Answer: C

Explanation:

The Notify & Suspend action for a resource monitor in Snowflake sends a notification to all account administrators who have notifications enabled and suspends all assigned warehouses. However, the suspension only occurs after all currently running statements in the warehouses have been completed¹. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 334

- (Topic 4)

Which metadata table will store the storage utilization information even for dropped tables?

- A. DATABASE_STORAGE_USAGE_HISTORY
- B. TABLE_STORAGE_METRICS
- C. STORAGE_DAILY_HISTORY
- D. STAGE STORAGE USAGE HISTORY

Answer: B

Explanation:

The TABLE_STORAGE_METRICS metadata table stores the storage utilization information, including for tables that have been dropped but are still incurring storage costs².

NEW QUESTION 335

- (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 337

- (Topic 5)

A user wants to add additional privileges to the system-defined roles for their virtual warehouse. How does Snowflake recommend they accomplish this?

- A. Grant the additional privileges to a custom role.
- B. Grant the additional privileges to the ACCOUNTADMIN role.
- C. Grant the additional privileges to the SYSADMIN role.
- D. Grant the additional privileges to the ORGADMIN role.

Answer: A

Explanation:

Snowflake recommends enhancing the granularity and management of privileges by creating and utilizing custom roles. When additional privileges are needed beyond those provided by the system-defined roles for a virtual warehouse or any other resource, these privileges should be granted to a custom role. This approach allows for more precise control over access rights and the ability to tailor permissions to the specific needs of different user groups or applications within the organization, while also maintaining the integrity and security model of system-defined roles.

References:

? Snowflake Documentation: Roles and Privileges

NEW QUESTION 342

- (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 343

- (Topic 5)

Which view can be used to determine if a table has frequent row updates or deletes?

- A. TABLES
- B. TABLE_STORAGE_METRICS
- C. STORAGE_DAILY_HISTORY
- D. STORAGE USAGE

Answer: B

Explanation:

The TABLE_STORAGE_METRICS view can be used to determine if a table has frequent row updates or deletes. This view provides detailed metrics on the storage utilization of tables within Snowflake, including metrics that reflect the impact of DML operations such as updates and deletes on table storage. For example, metrics related to the number of active and deleted rows can help identify tables that experience high levels of row modifications, indicating frequent updates or deletions.

References:

? Snowflake Documentation: TABLE_STORAGE_METRICS View

NEW QUESTION 344

- (Topic 5)

Which use case does the search optimization service support?

- A. Disjuncts (OR) in join predicates
- B. LIKE/ILIKE/RLIKE join predicates
- C. Join predicates on VARIANT columns
- D. Conjunctions (AND) of multiple equality predicates

Answer: D

Explanation:

The search optimization service in Snowflake supports use cases involving conjunctions (AND) of multiple equality predicates. This service enhances the performance of queries that include multiple equality conditions by utilizing search indexes to quickly filter data without scanning entire tables or partitions. It's particularly beneficial for improving the response times of complex queries that rely on specific data matching across multiple conditions.

References:

? Snowflake Documentation: Search Optimization Service

NEW QUESTION 346

- (Topic 5)

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

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

Answer: C

Explanation:

Snowflake bills virtual warehouse usage in one-minute increments, rounding up to the nearest minute for any partial minute of compute time used. If a virtual warehouse runs for 61 seconds and then, after being shut down, restarts and runs for an additional 30 seconds, the total time billed would be 120 seconds or 2 minutes. The first 61 seconds are rounded up to 2 minutes, and the subsequent 30 seconds are within a new minute, which is also rounded up to the nearest minute.

References:

? Snowflake Documentation: Virtual Warehouses Billing

NEW QUESTION 347

- (Topic 5)

What is it called when a customer managed key is combined with a Snowflake managed key to create a composite key for encryption?

- A. Hierarchical key model
- B. Client-side encryption
- C. Tri-secret secure encryption
- D. Key pair authentication

Answer: C

Explanation:

Tri-secret secure encryption is a security model employed by Snowflake that involves combining a customer-managed key with a Snowflake-managed key to create a composite key for encrypting data. This model enhances data security by requiring both the customer-managed key and the Snowflake-managed key to decrypt data, thus ensuring that neither party can access the data independently. It represents a balanced approach to key management, leveraging both customer control and Snowflake's managed services for robust data encryption.

References:

? Snowflake Documentation: Encryption and Key Management

NEW QUESTION 352

- (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 357

- (Topic 5)

Which data types optimally store semi-structured data? (Select TWO).

- A. ARRAY
- B. CHARACTER
- C. STRING
- D. VARCHAR
- E. VARIANT

Answer: AE

Explanation:

In Snowflake, semi-structured data is optimally stored using specific data types that are designed to handle the flexibility and complexity of such data. The **VARIANT** data type can store structured and semi-structured data types, including JSON, Avro, ORC, Parquet, or XML, in a single column. The **ARRAY** data type, on the other hand, is suitable for storing ordered sequences of elements, which can be particularly useful for semi-structured data types like JSON arrays. These data types provide the necessary flexibility to store and query semi-structured data efficiently in Snowflake.

References:

? Snowflake Documentation: Semi-structured Data Types

NEW QUESTION 361

- (Topic 5)

Which privilege is required to use the search optimization service in Snowflake?

- A. GRANT SEARCH OPTIMIZATION ON SCHEMA <schema_name> TO ROLE <role>
- B. GRANT SEARCH OPTIMIZATION ON DATABASE <database_name> TO ROLE <role>
- C. GRANT ADD SEARCH OPTIMIZATION ON SCHEMA <schema_name> TO ROLE <role>
- D. GRANT ADD SEARCH OPTIMIZATION ON DATABASE <database name> TO ROLE <role>

Answer: C

Explanation:

To utilize the search optimization service in Snowflake, the correct syntax for granting privileges to a role involves specific commands that include adding search optimization capabilities:

? Option C: GRANT ADD SEARCH OPTIMIZATION ON SCHEMA <schema_name>

TO ROLE <role>. This command grants the specified role the ability to implement search optimization at the schema level, which is essential for enhancing search capabilities within that schema.

Options A and B do not include the correct verb "ADD," which is necessary for this specific type of grant command in Snowflake. Option D incorrectly mentions the database level, as search optimization privileges are typically configured at the schema level, not the database level. References: Snowflake documentation on the use of GRANT statements for configuring search optimization.

NEW QUESTION 364

- (Topic 5)

Regardless of which notation is used, what are considerations for writing the column name and element names when traversing semi-structured data?

- A. The column name and element names are both case-sensitive.
- B. The column name and element names are both case-insensitive.
- C. The column name is case-sensitive but element names are case-insensitive.
- D. The column name is case-insensitive but element names are case-sensitive.

Answer: D

Explanation:

When querying semi-structured data in Snowflake, the behavior towards case sensitivity is distinct between column names and the names of elements within the semi-structured data. Column names follow the general SQL norm of being case-insensitive, meaning you can reference them in any case without affecting the query. However, element names within JSON, XML, or other semi-structured data are case-sensitive. This distinction is crucial for accurate data retrieval and manipulation in Snowflake, especially when working with JSON objects where the case of keys can significantly alter the outcome of queries.

References:

? Snowflake Documentation: Querying Semi-structured Data

NEW QUESTION 366

- (Topic 5)

What type of function returns one value for each Invocation?

- A. Aggregate
- B. Scalar
- C. Table
- D. Window

Answer: B

Explanation:

Scalar functions in Snowflake (and SQL in general) are designed to return a single value for each invocation. They operate on a single value and return a single result, making them suitable for a wide range of data transformations and calculations within queries.

References:

? Snowflake Documentation: Functions

NEW QUESTION 369

- (Topic 5)

Which role has the ability to create a share from a shared database by default?

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

Answer: A

Explanation:

By default, the ACCOUNTADMIN role in Snowflake has the ability to create a share from a shared database. This role has the highest level of access within a Snowflake account, including the management of all aspects of the account, such as users, roles, warehouses, and databases, as well as the creation and management of shares for secure data sharing with other Snowflake accounts.

References:

? Snowflake Documentation: Roles

NEW QUESTION 371

- (Topic 5)

How are network policies defined in Snowflake?

- A. They are a set of rules that define the network routes within Snowflake.
- B. They are a set of rules that dictate how Snowflake accounts can be used between multiple users.
- C. They are a set of rules that define how data can be transferred between different Snowflake accounts within an organization.
- D. They are a set of rules that control access to Snowflake accounts by specifying the IP addresses or ranges of IP addresses that are allowed to connect to Snowflake.

Answer: D

Explanation:

Network policies in Snowflake are defined as a set of rules that manage the network-level access to Snowflake accounts. These rules specify which IP addresses or IP ranges are permitted to connect to Snowflake, enhancing the security of Snowflake accounts by preventing unauthorized access. Network policies are an essential aspect of Snowflake's security model, allowing administrators to enforce access controls based on network locations.

References:

? Snowflake Documentation: Network Policies

NEW QUESTION 373

- (Topic 5)

What does the worksheet and database explorer feature in Snowsight allow users to do?

- A. Add or remove users from a worksheet.
- B. Move a worksheet to a folder or a dashboard.
- C. Combine multiple worksheets into a single worksheet.
- D. Tag frequently accessed worksheets for ease of access.

Answer: D

Explanation:

The worksheet and database explorer feature in Snowsight allows users to tag frequently accessed worksheets for ease of access. This functionality helps users organize and quickly navigate to the worksheets they use most often, enhancing productivity and streamlining the data exploration and analysis process within Snowsight, Snowflake's web-based query and visualization interface.

References:

? Snowflake Documentation: Snowsight (UI for Snowflake)

NEW QUESTION 374

- (Topic 5)

How long is a query visible in the Query History page in the Snowflake Web Interface (UI)?

- A. 60 minutes
- B. 24 hours
- C. 14 days
- D. 30 days

Answer: C

Explanation:

In the Snowflake Web Interface (UI), the Query History page displays the history of queries executed in Snowflake for up to 14 days. This allows users to review and analyze their query performance, troubleshoot issues, and understand their query patterns over a two-week period. The Query History page is a critical tool for monitoring and optimizing the use of Snowflake.

References:

? Snowflake Documentation: Using the Web Interface

NEW QUESTION 379

- (Topic 5)

Which SQL command can be used to verify the privileges that are granted to a role?

- A. SHOW GRANTS ON ROLE <Role Name>
- B. SHOW ROLES <Role Name>
- C. SHOW GRANTS TO ROLE <Role Name>
- D. SHOW GRANTS FOR ROLE <Role Name>

Answer: C

Explanation:

To verify the privileges that have been granted to a specific role in Snowflake, the correct SQL command is SHOW GRANTS TO ROLE <Role Name>. This command lists all the privileges granted to the specified role, including access to schemas, tables, and other database objects. This is a useful command for administrators and users with sufficient privileges to audit and manage role permissions within the Snowflake environment.

References:

? Snowflake Documentation: SHOW GRANTS

NEW QUESTION 384

- (Topic 5)

What are characteristics of transient tables in Snowflake? (Select TWO).

- A. Transient tables have a Fail-safe period of 7 days.
- B. Transient tables can be cloned to permanent tables.
- C. Transient tables persist until they are explicitly dropped.
- D. Transient tables can be altered to make them permanent tables.
- E. Transient tables have Time Travel retention periods of 0 or 1 day.

Answer: BC

Explanation:

Transient tables in Snowflake are designed for temporary or intermediate workloads with the following characteristics:

? B. Transient tables can be cloned to permanent tables: This feature allows users to create copies of transient tables for permanent use, providing flexibility in managing data lifecycles.

? C. Transient tables persist until they are explicitly dropped: Unlike temporary tables that exist for the duration of a session, transient tables remain in the database until explicitly removed by a user, offering more durability for short-term data storage needs.

References:

? Snowflake Documentation: Transient Tables

NEW QUESTION 388

- (Topic 5)

Which command can be used to list all the file formats for which a user has access privileges?

- A. LIST
- B. ALTER FILE FORMAT
- C. DESCRIBE FILE FORMAT
- D. SHOW FILE FORMATS

Answer: D

Explanation:

The command to list all the file formats for which a user has access privileges in Snowflake is SHOW FILE FORMATS. This command provides a list of all file formats defined in the user's current session or specified database/schema, along with details such as the name, type, and creation time of each file format. It is a valuable tool for users to understand and manage the file formats available for data loading and unloading operations.

References:

? Snowflake Documentation: SHOW FILE FORMATS

NEW QUESTION 390

- (Topic 5)

What is the MAXIMUM number of clusters that can be provisioned with a multi-cluster virtual warehouse?

- A. 1
- B. 5
- C. 10
- D. 100

Answer: C

Explanation:

In Snowflake, the maximum number of clusters that can be provisioned within a multi-cluster virtual warehouse is 10. This allows for significant scalability and performance management by enabling Snowflake to handle varying levels of query load by adjusting the number of active clusters within the warehouse. References: Snowflake documentation on virtual warehouses, particularly the scalability options available in multi-cluster configurations.

NEW QUESTION 392

- (Topic 5)

What is the Fail-safe retention period for transient and temporary tables?

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

Answer: A

Explanation:

The Fail-safe retention period for transient and temporary tables in Snowflake is 0 days. Fail-safe is a feature designed to protect data against accidental loss or deletion by retaining historical data for a period after its Time Travel retention period expires. However, transient and temporary tables, which are designed for temporary or short-term storage and operations, do not have a Fail-safe period. Once the data is deleted or the table is dropped, it cannot be recovered.

References:

? Snowflake Documentation: Understanding Fail-safe

NEW QUESTION 396

- (Topic 5)

User1, who has the SYSADMIN role, executed a query on Snowsight. User2, who is in the same Snowflake account, wants to view the result set of the query executed by User1 using the Snowsight query history.

What will happen if User2 tries to access the query history?

- A. If User2 has the sysadmin role they will be able to see the results.
- B. If User2 has the securityadmin role they will be able to see the results.
- C. If User2 has the ACCOUNTADMIN role they will be able to see the results.
- D. User2 will be unable to view the result set of the query executed by User1.

Answer: C

Explanation:

In Snowflake, the query history and the results of queries executed by a user are accessible based on the roles and permissions. If User1 executed a query with the SYSADMIN role, User2 would be able to view the result set of that query executed by User1 only if User2 has the ACCOUNTADMIN role. The ACCOUNTADMIN role has the broadest set of privileges, including the ability to access all aspects of the account's operation, data, and query history, thus enabling User2 to view the results of queries executed by other users.

References:

? Snowflake Documentation: Understanding Snowflake Roles

NEW QUESTION 398

- (Topic 5)

Which function should be used to insert JSON format string data into a VARIANT field?

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

Answer: C

Explanation:

To insert JSON formatted string data into a VARIANT field in Snowflake, the correct function to use is PARSE_JSON. The PARSE_JSON function is specifically designed to interpret a JSON formatted string and convert it into a VARIANT type, which is Snowflake's flexible format for handling semi-structured data like JSON, XML, and Avro. This function is essential for loading and querying JSON data within Snowflake, allowing users to store and manage JSON data efficiently while preserving its structure for querying purposes. This function's usage and capabilities are detailed in the Snowflake documentation, providing users with guidance on how to handle semi-structured data effectively within their Snowflake environments.

References:

? Snowflake Documentation: PARSE_JSON

NEW QUESTION 399

- (Topic 5)

Which command removes a role from another role or a user in Snowflak?

- A. ALTER ROLE
- B. REVOKE ROLE
- C. USE ROLE
- D. USE SECONDARY ROLES

Answer: B

Explanation:

The **REVOKE ROLE** command is used to remove a role from another role or a user in Snowflake. This command is part of Snowflake's role-based access control system, allowing administrators to manage permissions and access to database objects efficiently by adding or removing roles from users or other roles.

References:

? Snowflake Documentation: REVOKE ROLE

NEW QUESTION 403

- (Topic 5)

What criteria does Snowflake use to determine the current role when initiating a session? (Select TWO).

- A. If a role was specified as part of the connection and that role has been granted to the Snowflake user, the specified role becomes the current role.
- B. If no role was specified as part of the connection and a default role has been defined for the Snowflake user, that role becomes the current role.
- C. If no role was specified as part of the connection and a default role has not been set for the Snowflake user, the session will not be initiated and the log in will fail.
- D. If a role was specified as part of the connection and that role has not been granted to the Snowflake user, it will be ignored and the default role will become the current role.
- E. If a role was specified as part of the connection and that role has not been granted to the Snowflake user, the role is automatically granted and it becomes the current role.

Answer: AB

Explanation:

When initiating a session in Snowflake, the system determines the current role based on the user's connection details and role assignments. If a user specifies a role during the connection, and that role is already granted to them, Snowflake sets it as the current role for the session. Alternatively, if no role is specified during the connection, but the user has a default role assigned, Snowflake will use this default role as the current session role. These mechanisms ensure that users operate within their permissions, enhancing security and governance within Snowflake environments.

References:

? Snowflake Documentation: Understanding Roles

NEW QUESTION 405

- (Topic 5)

What will happen if a Snowflake user increases the size of a suspended virtual warehouse?

- A. The provisioning of new compute resources for the warehouse will begin immediately.
- B. The warehouse will remain suspended but new resources will be added to the query acceleration service.
- C. The provisioning of additional compute resources will be in effect when the warehouse is next resumed.
- D. The warehouse will resume immediately and start to share the compute load with other running virtual warehouses.

Answer: C

Explanation:

When a Snowflake user increases the size of a suspended virtual warehouse, the changes to compute resources are queued but do not take immediate effect. The provisioning of additional compute resources occurs only when the warehouse is resumed. This ensures that resources are allocated efficiently, aligning with Snowflake's commitment to cost-effective and on-demand scalability.

References:

? Snowflake Documentation: Virtual Warehouses

NEW QUESTION 408

- (Topic 5)

What should be used when creating a CSV file format where the columns are wrapped by single quotes or double quotes?

- A. **BINARY_FORMAT**
- B. **ESCAPE_UNENCLOSED_FIELD**
- C. **FIELD_OPTIONALLY_ENCLOSED_BY**
- D. **SKIP_BYTE_ORDER MARK**

Answer: C

Explanation:

When creating a CSV file format in Snowflake and the requirement is to wrap columns by single quotes or double quotes, the **FIELD_OPTIONALLY_ENCLOSED_BY** parameter should be used in the file format specification. This parameter allows you to define a character (either a single quote or a double quote) that can optionally enclose each field in the CSV file, providing flexibility in handling fields that contain special characters or delimiters as part of their data.

References:

? Snowflake Documentation: CSV File Format

NEW QUESTION 410

- (Topic 5)

What is the MINIMUM permission needed to access a file URL from an external stage?

- A. MODIFY
- B. READ
- C. SELECT
- D. USAGE

Answer: D

Explanation:

To access a file URL from an external stage in Snowflake, the minimum permission required is USAGE on the stage object. USAGE permission allows a user to reference the stage in SQL commands, necessary for actions like listing files or loading data from the stage, but does not permit the user to alter or drop the stage.

References:

? Snowflake Documentation: Access Control

NEW QUESTION 415

- (Topic 5)

When floating-point number columns are unloaded to CSV or JSON files, Snowflake truncates the values to approximately what?

- A. (12,2)
- B. (10,4)
- C. (14,8)
- D. (15,9)

Answer: D

Explanation:

When unloading floating-point number columns to CSV or JSON files, Snowflake truncates the values to approximately 15 significant digits with 9 digits following the decimal point, which can be represented as (15,9). This ensures a balance between accuracy and efficiency in representing floating-point numbers in text-based formats, which is essential for data interchange and processing applications that consume these files. References:

? Snowflake Documentation: Data Unloading Considerations

NEW QUESTION 419

- (Topic 5)

A Snowflake user is writing a User-Defined Function (UDF) that includes some unqualified object names.

How will those object names be resolved during execution?

- A. Snowflake will resolve them according to the SEARCH_PATH parameter.
- B. Snowflake will only check the schema the UDF belongs to.
- C. Snowflake will first check the current schema, and then the schema the previous query used.
- D. Snowflake will first check the current schema, and then the PUBLIC schema of the current database.

Answer: D

Explanation:

? Object Name Resolution: When unqualified object names (e.g., table name without schema) are used in a UDF, Snowflake follows a specific hierarchy to resolve them. Here's the order:

? Note: The SEARCH_PATH parameter influences object resolution for queries, not within UDFs.

References:

? Snowflake Documentation (Object Naming

Resolution): <https://docs.snowflake.com/en/sql-reference/name-resolution.html>

NEW QUESTION 424

- (Topic 6)

A Snowflake user is trying to load a 125 GB file using SnowSQL. The file continues to load for almost an entire day. What will happen at the 24-hour mark?

- A. The file will continue to load until all contents are loaded.
- B. The file will stop loading and all data up to that point will be committed.
- C. The file loading could be aborted without any portion of the file being committed.
- D. The file's number of allowable hours to load can be programmatically controlled to load easily into Snowflake.

Answer: C

Explanation:

When attempting to load large files, such as a 125 GB file, into Snowflake using SnowSQL, the process might encounter limitations related to the maximum execution time for queries or data loading operations. If the loading process exceeds this time limit (typically around 24 hours), it could be aborted without committing any part of the file to the database. This behavior is designed to prevent indefinite resource consumption and to maintain system stability, emphasizing the need for optimizing data load operations, possibly through file segmentation or parallel loading strategies. References: Snowflake Documentation on Data Loading Considerations

NEW QUESTION 427

- (Topic 6)

Which function can be used with the copy into <location> statement to convert rows from a relational table to a single variant column, and to unload rows into a JSON file?

- A. FLATTEN
- B. OBJECT_AS
- C. OBJECT_CONSTRUCT
- D. TO VARIANT

Answer: D

Explanation:

The correct function to use with the COPY INTO <location> statement to convert rows from a relational table into a single variant column and to unload rows into a JSON file is TO VARIANT. The TO VARIANT function is used to explicitly convert a value of any supported data type into a VARIANT data type. This is particularly useful when needing to aggregate multiple columns or complex data structures into a single JSON-formatted string, which can then be unloaded into a file. In the context of unloading data, the COPY INTO <location> statement combined with TO VARIANT enables the conversion of structured data from Snowflake tables into a semi-structured VARIANT format, typically JSON, which can then be efficiently exported and stored. This approach is often utilized for data integration scenarios, backups, or when data needs to be shared in a format that is easily consumed by various applications or services that support JSON.

References:

? Snowflake Documentation on Data Unloading: Unloading Data

? Snowflake Documentation on VARIANT Data Type: Working with JSON

NEW QUESTION 430

- (Topic 6)

Which object type is granted permissions for reading a table?

- A. User
- B. Role
- C. Attribute
- D. Schema

Answer: B

Explanation:

In Snowflake, permissions for accessing database objects, including tables, are not granted directly to users but rather to roles. A role encapsulates a collection of privileges on various Snowflake objects. Users are then granted roles, and through those roles, they inherit the permissions necessary to read a table or perform other actions. This approach adheres to the principle of least privilege, allowing for granular control over database access and simplifying the management of user permissions.

Reference: Snowflake's official documentation on access control introduces the concept of

roles and how they are used to manage permissions: <https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#roles>

NEW QUESTION 435

- (Topic 6)

Who can create network policies within Snowflake? (Select TWO).

- A. SYSADMIN only
- B. ORCADMIN only
- C. SECURITYADMIN or higher roles
- D. A role with the CREATE NETWORK POLICY privilege
- E. A role with the CREATE SECURITY INTEGRATION privilege

Answer: CD

Explanation:

In Snowflake, network policies define the allowed IP address ranges from which users can connect to Snowflake, enhancing security by restricting access based on network location. The creation and management of network policies require sufficient privileges. Specifically, a user with the SECURITYADMIN role or any role with higher privileges, such as ACCOUNTADMIN, can create network policies. Additionally, a custom role can be granted the CREATE NETWORK POLICY privilege, enabling users assigned to that role to also create network policies. This approach allows for flexible and secure management of network access to Snowflake.

References: Snowflake Documentation on Network Policies

NEW QUESTION 439

- (Topic 6)

Which actions can be performed using a resource monitor in Snowflake? (Select TWO).

- A. Monitor the performance of individual queries in real-time
- B. Automatically allocate more storage space to a virtual warehouse
- C. Modify the queries being executed within a virtual warehouse.
- D. Suspend a virtual warehouse when its credit usage reaches a defined limit.
- E. Trigger a notification to account administrators when credit usage reaches a specified threshold

Answer: DE

Explanation:

Resource monitors in Snowflake can perform actions such as suspending a virtual warehouse when its credit usage reaches a defined limit and triggering a notification to account administrators when credit usage reaches a specified threshold. These actions help manage and control resource usage and costs within Snowflake.

References:

? Snowflake Documentation: Resource Monitors

NEW QUESTION 440

- (Topic 6)

Which MINIMUM set of privileges is required to temporarily bypass an active network policy by configuring the user object property MINS_TO_BYPASS_NETWORK_POLICY?

- A. Only while in the ACCOUNTADMIN role
- B. Only while in the securityadmin role
- C. Only the role with the ownership privilege on the network policy
- D. Only Snowflake Support can set the value for this object property

Answer: A

Explanation:

To temporarily bypass an active network policy by configuring the user object property MINS_TO_BYPASS_NETWORK_POLICY, the minimum set of privileges required is having the ACCOUNTADMIN role. This role has the necessary privileges to make such changes, including modifying user properties that affect network policies.

References:

? Snowflake Documentation: Network Policy Management

NEW QUESTION 445

- (Topic 6)

Which security models are used in Snowflake to manage access control? (Select TWO).

- A. Discretionary Access Control (DAC)
- B. Identity Access Management (IAM)
- C. Mandatory Access Control (MAC)
- D. Role-Based Access Control (RBAC)
- E. Security Assertion Markup Language (SAML)

Answer: AD

Explanation:

Snowflake uses both Discretionary Access Control (DAC) and Role-Based Access Control (RBAC) to manage access control. DAC allows object owners to grant access privileges to other users. RBAC assigns permissions to roles, and roles are then granted to users, making it easier to manage permissions based on user roles within the organization.

References:

? Snowflake Documentation: Access Control in Snowflake

NEW QUESTION 447

- (Topic 6)

Which statistics on a Query Profile reflect the efficiency of the query pruning? (Select TWO).

- A. Partitions scanned
- B. Partitions total
- C. Bytes spilled
- D. Bytes scanned
- E. Bytes written

Answer: AD

Explanation:

In a Snowflake Query Profile, the statistics "Partitions scanned" and "Bytes scanned" reflect the efficiency of query pruning. Query pruning refers to the ability of the query engine to skip unnecessary data, thereby reducing the amount of data that needs to be processed. Efficient pruning results in fewer partitions and bytes being scanned, improving query performance.

References:

? Snowflake Documentation: Understanding Query Profiles Top of Form Bottom of Form

NEW QUESTION 450

- (Topic 6)

What optional properties can a Snowflake user set when creating a virtual warehouse? (Select TWO).

- A. Auto-suspend
- B. Cache size
- C. Default role
- D. Resource monitor
- E. Storage size

Answer: AD

Explanation:

When creating a virtual warehouse in Snowflake, users have the option to set several properties to manage its behavior and resource usage. Two of these optional properties are Auto-suspend and Resource monitor.

? Auto-suspend: This property defines the period of inactivity after which the warehouse will automatically suspend. This helps in managing costs by stopping the warehouse when it is not in use.

```
CREATE WAREHOUSE my_warehouse WITH WAREHOUSE_SIZE = 'XSMALL'
```

```
AUTO_SUSPEND = 300; -- Auto-suspend after 5 minutes of inactivity
```

? Resource monitor: Users can assign a resource monitor to a warehouse to control and limit the amount of credit usage. Resource monitors help in setting quotas and alerts for warehouse usage.

```
CREATE WAREHOUSE my_warehouse WITH WAREHOUSE_SIZE = 'XSMALL'
```

```
RESOURCE_MONITOR = 'my_resource_monitor';
```

References:

? Snowflake Documentation: Creating Warehouses

? Snowflake Documentation: Resource Monitors

NEW QUESTION 452

- (Topic 6)

Which service or tool is a Command Line Interface (CLI) client used for connecting to Snowflake to execute SQL queries?

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

Answer: D

Explanation:

SnowSQL is the Command Line Interface (CLI) client provided by Snowflake for executing SQL queries and performing various tasks. It allows users to connect to their Snowflake accounts and interact with the Snowflake data warehouse.

? Installation: SnowSQL can be downloaded and installed on various operating systems.

? Configuration: Users need to configure SnowSQL with their Snowflake account credentials.

? Usage: Once configured, users can run SQL queries, manage data, and perform administrative tasks through the CLI.

References:

? Snowflake Documentation: SnowSQL

? Snowflake Documentation: Installing SnowSQL

NEW QUESTION 457

- (Topic 6)

What Snowflake objects can contain custom application logic written in JavaScript? (Select TWO)

- A. Stored procedures
- B. Stages
- C. Tasks
- D. Views
- E. User-Defined Functions (UDFs)

Answer: AE

Explanation:

Snowflake allows users to write custom application logic in JavaScript for two types of objects: Stored Procedures and User-Defined Functions (UDFs).

? Stored Procedures: Snowflake stored procedures can be written in JavaScript to encapsulate complex business logic and procedural operations. CREATE OR REPLACE PROCEDURE my_procedure()

RETURNS STRING LANGUAGE JAVASCRIPT EXECUTE AS CALLER AS

\$\$

// JavaScript logic here

\$\$;

? User-Defined Functions (UDFs): Snowflake UDFs can be written in JavaScript to perform custom calculations or operations on data.

CREATE OR REPLACE FUNCTION my_function(x FLOAT) RETURNS FLOAT

LANGUAGE JAVASCRIPT AS

\$\$

return x * 2;

\$\$;

References:

? Snowflake Documentation: Stored Procedures

? Snowflake Documentation: User-Defined Functions (UDFs)

NEW QUESTION 458

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

COF-C02 Practice Exam Features:

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

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