

Snowflake.COF-C03.v2026-09-29.q478

Exam Code:	COF-C03
Exam Name:	SnowPro® Core Certification 2026 Exam (COF-C03)
Certification Provider:	Snowflake
Free Question Number:	478
Version:	v2026-09-29
# of views:	102
# of Questions views:	4793
https://www.freecram.net/torrent/Snowflake.COF-C03.v2026-09-29.q478.html	

NEW QUESTION: 1

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 2

Which Snowflake edition offers the highest level of security for organizations that have the strictest requirements?

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

Answer: ([SHOW ANSWER](#))

The Virtual Private Snowflake (VPS) edition offers the highest level of security for organizations with the strictest security requirements. This edition provides a dedicated and isolated instance of Snowflake, including enhanced security features and compliance certifications to meet the needs of highly regulated industries or any organization requiring the utmost in data protection and privacy.

References:

Snowflake Documentation: Snowflake Editions

NEW QUESTION: 3

How does Snowflake describe its unique architecture?

- A. A single-cluster shared data architecture using a central data repository and massively parallel processing (MPP)
- B. A multi-duster shared nothing architecture using a soloed data repository and massively parallel processing (MPP)
- C. A single-cluster shared nothing architecture using a sliced data repository and symmetric multiprocessing (SMP)
- D. A multi-cluster shared nothing architecture using a siloed data repository and symmetric multiprocessing (SMP)

Answer: ([SHOW ANSWER](#))

Snowflake's unique architecture is described as a multi-cluster, shared data architecture that leverages massively parallel processing (MPP). This architecture separates compute and storage resources, enabling Snowflake to scale them independently. It does not use a single cluster or rely solely on symmetric multiprocessing (SMP); rather, it uses a combination of shared-nothing architecture for compute clusters (virtual warehouses) and a centralized storage layer for data, optimizing for both performance and scalability.

References:

Snowflake Documentation: Snowflake Architecture Overview

NEW QUESTION: 4

When creating a custom role that will be allowed to create warehouses and databases, which system-defined role does Snowflake recommend also be assigned to the custom role?

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 5

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

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

Answer: ([SHOW ANSWER](#))

Snowflake supports several semi-structured data formats for loading data. The supported formats include JSON, Avro, and Parquet¹². These formats allow for efficient storage and querying of data that does not conform to a traditional relational database schema.

NEW QUESTION: 6

A user needs to MINIMIZE the cost of large tables that are used to store transitory data. The data does not need to be protected against failures, because the data can be reconstructed outside of Snowflake.

What table type should be used?

- A. Permanent
- B. Transient

C. Temporary

D. External

Answer: ([SHOW ANSWER](#))

For minimizing the cost of large tables that are used to store transitory data, which does not need to be protected against failures because it can be reconstructed outside of Snowflake, the best table type to use is Transient. Transient tables in Snowflake are designed for temporary or transitory data storage and offer reduced storage costs compared to permanent tables. However, unlike temporary tables, they persist across sessions until explicitly dropped.

Why Transient Tables: Transient tables provide a cost-effective solution for storing data that is temporary but needs to be available longer than a single session. They have lower data storage costs because Snowflake does not maintain historical data (Time Travel) for as long as it does for permanent tables.

Creating a Transient Table:

To create a transient table, use the `TRANSIENT` keyword in the `CREATE TABLE` statement:

```
CREATE TRANSIENT TABLE my_transient_table (...);
```

Use Case Considerations: Transient tables are ideal for scenarios where the data is not critical, can be easily recreated, and where cost optimization is a priority. They are suitable for development, testing, or staging environments where data longevity is not a concern.

Reference: For more details on transient tables and their usage scenarios, refer to the Snowflake documentation on table types: <https://docs.snowflake.com/en/sql-reference/sql/create-table.html#table-types>

NEW QUESTION: 7

What happens when a database is cloned?

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

Answer: ([SHOW ANSWER](#))

When a database is cloned in Snowflake, it does not retain any privileges that were granted on the source object. The clone will need to have privileges reassigned as necessary for users to access it. References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 8

Which technical needs require Enterprise Edition? Select TWO.

- A. The need to use network policies to limit site access by user IP address.
- B. The need for access to private connectivity.
- C. The need to use multi-cluster virtual warehouses for high concurrency.
- D. The need to use column-level security.
- E. The need to use Tri-Secret Secure.

Answer: ([SHOW ANSWER](#))

Enterprise Edition includes features such as multi-cluster warehouses and column-level security features.

Multi-cluster warehouses are used for high concurrency, and column-level security includes features such as masking policies.

Why the others are incorrect:

Network policies are available below Enterprise. Private connectivity and Tri-Secret Secure require Business Critical or higher, not merely Enterprise.

Reference: Snowflake Documentation - Snowflake editions; Multi-cluster warehouses; Column-level security.

NEW QUESTION: 9

Which Snowflake object contains all the information required to share a database?

- A. Private listing
- B. Secure view
- C. Sequence
- D. Share

Answer: ([SHOW ANSWER](#))

In Snowflake, a Share is the object that contains all the information required to share a database with other Snowflake accounts. Shares are used to securely share data stored in Snowflake tables and views, enabling data providers to grant data consumers access to their datasets without duplicating data. When a database is shared, it can include one or more schemas, and each schema can contain tables, views, or both.

References:

Snowflake Documentation on Shares: Shares

NEW QUESTION: 10

In Snowflake's data security framework, how does column-level security contribute to the protection of sensitive information? (Select TWO).

- A. Implementation of column-level security will optimize query performance.
- B. Column-level security supports encryption of the entire database.
- C. Column-level security ensures that only the table owner can access the data.
- D. Column-level security limits access to specific columns within a table based on user privileges
- E. Column-level security allows the application of a masking policy to a column within a table or view.

Answer: ([SHOW ANSWER](#))

Column-level security in Snowflake enhances data protection by restricting access and applying masking policies to sensitive data at the column level.

Limiting Access Based on User Privileges:

Column-level security allows administrators to define which users or roles have access to specific columns within a table.

This ensures that sensitive data is only accessible to authorized personnel, thereby reducing the risk of data breaches.

Application of Masking Policies:

Masking policies can be applied to columns to obfuscate sensitive data.

For example, credit card numbers can be masked to show only the last four digits, protecting the full number from being exposed.

Snowflake Documentation: Column-Level Security

Snowflake Documentation: Dynamic Data Masking

NEW QUESTION: 11

Which commands are supported by the query acceleration service? (Select TWO).

- A. SELECT
- B. INSERT
- C. MERGE
- D. DELETE
- E. UPDATE

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 12

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: ([SHOW ANSWER](#))

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 meaning.

References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 13

Which Snowflake SQL statement would be used to determine which users and roles have access to a role called MY_ROLE?

- A. SHOW GRANTS OF ROLE MY_ROLE
- B. SHOW GRANTS TO ROLE MY_ROLE
- C. SHOW GRANTS FOR ROLE MY_ROLE
- D. SHOW GRANTS ON ROLE MY_ROLE

Answer: ([SHOW ANSWER](#))

The SQL statement SHOW GRANTS TO ROLE MY_ROLE is used to determine which users and roles have access to a role called MY_ROLE. This statement lists all the privileges granted to the role, including which roles and users can assume MY_ROLE. References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 14

A user wants to access stored in a stage without authenticating into Snowflake.

Which type of URL should be used?

- A. File URL
- B. Staged URL
- C. Scoped URL

D. Pre-signed URL

Answer: D ([LEAVE A REPLY](#))

A pre-signed URL provides a way to access objects stored in a stage without requiring authentication to Snowflake. This URL contains all the necessary information for access control, including an expiration time, and it's signed by the service 's credentials. This mechanism allows users or applications to access specific files stored in a cloud storage location (stage) directly, bypassing the need to authenticate into Snowflake, ensuring secure and temporary access to the data.

References: Snowflake Documentation on Stages and Data Loading

NEW QUESTION: 15

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: ([SHOW ANSWER](#))

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: 16

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

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

Answer: ([SHOW ANSWER](#))

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

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:

Code: **freecram**)

NEW QUESTION: 17

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: ([SHOW ANSWER](#))

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: 18

What is a User-Defined Table Function, or UDTF, used for in Snowflake?

- A. To return a tabular value for each input row.
- B. To process input rows and return the results as a pandas series.
- C. To return one output row as a single column or value for each input row.
- D. To process input rows as pandas DataFrames and return tabular results.

Answer: ([SHOW ANSWER](#))

The correct answer is A. To return a tabular value for each input row.

A User-Defined Table Function, or UDTF, returns a set of rows instead of a single scalar value. This is different from a scalar UDF, which returns one value for each input row.

Why A is correct:

A UDTF produces tabular output. It can return multiple rows and multiple columns, depending on how the function is defined.

Example concept:

```
SELECT *  
FROM TABLE(my_udtf(input_column));
```

This syntax shows that the function result is treated like a table.

Why the other options are incorrect:

- B). Returning a pandas Series is associated with vectorized Python UDF behavior, not the general definition of a UDTF.
- C). Returning one output row as a single column or value describes a scalar UDF, not a table function.
- D). Processing input rows as pandas DataFrames is associated with Snowpark Python and vectorized processing patterns, but it is not the basic Snowflake definition of a UDTF.

Official Snowflake documentation reference:

Snowflake documentation defines a UDTF as a user-defined function that returns tabular results. It can return multiple rows for each input row.

Reference: Snowflake Documentation - User-defined table functions; Snowflake Documentation - User-defined functions; SnowPro Core Study Guide - SQL and Snowflake Objects.

NEW QUESTION: 19

What does creating a resource monitor for a virtual warehouse allow a user to do?

- A. Schedule when the warehouse starts and stops.
- B. Track and control the credits consumed by the warehouse.
- C. Monitor the real-time performance of all queries running on the warehouse.
- D. Automatically adjust the settings of the warehouse based on the workload.

Answer: ([SHOW ANSWER](#))

Resource monitors are used to track and control credit usage. They can trigger notifications or suspend compute resources when thresholds are exceeded.

Reference:

Snowflake Docs: Resource Monitors

NEW QUESTION: 20

A company stores Personal Identifiable Information (PII) in Snowflake across multiple countries, including the United States, Canada, and the United Kingdom. How can Snowflake be used to automatically identify data, in order to support data governance requirements for each country?

- A. Use semantic category classification.
- B. Create a custom User-Defined Function (UDF).
- C. Implement a third-party service to protect the data.
- D. Tag all columns directly that contain PII.

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 21

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

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

Answer: ([SHOW ANSWER](#))

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: 22

For Directory tables, what stage allows for automatic refreshing of metadata?

- A. User stage
- B. Table stage
- C. Named internal stage

D. Named external stage

Answer: ([SHOW ANSWER](#))

For directory tables, a named external stage allows for the automatic refreshing of metadata. This capability is particularly useful when dealing with files stored on external storage services (like Amazon S3, Google Cloud Storage, or Azure Blob Storage) and accessed through Snowflake. The external stage references these files, and the directory table 's metadata can be automatically updated to reflect changes in the underlying files.

References:

Snowflake Documentation: External Stages

NEW QUESTION: 23

Which COPY INTO < table > options can be used to load data from a staged file when the metadata in the file has expired? (Select TWO)

- A.** LOAD_EXPIRED_FILES
- B.** FORCE
- C.** ENFORCE_LENGTH
- D.** PURGE
- E.** LOAD_UNCERTAIN_FILES

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 24

Which steps will reduce memory spillage when running a very large query? (Select TWO)

- A.** Enable auto-scale mode on the virtual warehouse.
- B.** Increase the network bandwidth.
- C.** Process the data in smaller batches.
- D.** Add a DISTINCT clause to the query.
- E.** Increase the size of the virtual warehouse.

Answer: ([SHOW ANSWER](#))

* Processing data in smaller batches reduces resource pressure during execution.

* Increasing warehouse size provides more memory and CPU, which helps reduce memory spillage , where intermediate data is written to disk.

Reference:

Snowflake Docs: Query Performance Tuning

NEW QUESTION: 25

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 26

A table gets updated throughout the day. A user wants to automate the process of table changes by capturing only the changed rows every 20 minutes and applying transformations without manual intervention. How can this be accomplished?

- A. Create a stream and run a task only when there are changes.
- B. Create a task that triggers every 20 minutes.
- C. Create a stream and schedule a task every 20 minutes.
- D. Use a dynamic table to materialize changes and run a task for the business transformations every 20 minutes.

Answer: ([SHOW ANSWER](#))

The correct answer is C. Create a stream and schedule a task every 20 minutes.

A stream is used to capture change data capture, or CDC, information for a table. It records inserts, updates, and deletes made to the source table. A task is used to automate SQL execution on a schedule. Together, streams and tasks are commonly used to automate incremental data processing in Snowflake.

Why C is correct:

The requirement has two parts:

Capture only changed rows.

Run the transformation automatically every 20 minutes.

A stream satisfies the first requirement because it tracks changed rows. A scheduled task satisfies the second requirement because it can run SQL logic at a defined interval, such as every 20 minutes.

Example concept:

```
CREATE STREAM my_stream
ON TABLE source_table;
CREATE TASK my_task
WAREHOUSE = my_warehouse
SCHEDULE = ' 20 MINUTE '
AS
INSERT INTO target_table
SELECT ...
FROM my_stream;
```

Why the other options are incorrect:

A). This is partially correct, but it does not clearly satisfy the requirement to run every 20 minutes. A task can include a `WHEN SYSTEM$STREAM_HAS_DATA(...)` condition, but the question specifically asks for processing every 20 minutes.

B). A task alone can run every 20 minutes, but it does not capture only changed rows.

D). Dynamic tables are used for declarative automated data transformation, but the question specifically asks to capture only changed rows from a frequently updated table, which is the purpose of a stream.

Official Snowflake documentation reference:

Snowflake documentation describes streams as objects that record data manipulation language, or DML, changes made to tables. Snowflake tasks execute SQL statements on a schedule or as part of a task graph.

Reference: Snowflake Documentation - Streams; Snowflake Documentation - Introduction to tasks; SnowPro Core Study Guide - SQL and Snowflake Objects.

NEW QUESTION: 27

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 ([LEAVE A REPLY](#))

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: 28

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: ([SHOW ANSWER](#)**)**

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: 29

A multi-cluster virtual warehouse is running in Auto-scale mode. What happens if the number of concurrent queries exceeds the capacity of the currently running cluster?

- A. Queued queries will be dropped.
- B. The warehouse will shut down.
- C. A new cluster will be started.
- D. Queries will be redirected to another warehouse.

Answer: ([SHOW ANSWER](#))

The correct answer is C. A new cluster will be started.

A multi-cluster warehouse in Auto-scale mode can automatically start additional clusters when query concurrency increases. This allows Snowflake to handle more concurrent users and queries without forcing all queries to wait on a single cluster.

Why C is correct:

When the current cluster does not have enough capacity to handle the number of concurrent queries, Snowflake can start another cluster, up to the configured maximum cluster count. This helps reduce queuing and improves concurrency.

Why the other options are incorrect:

A). Queued queries are not dropped simply because concurrency exceeds the current cluster's capacity.

B). The warehouse does not shut down because of high concurrency.

D). Queries are not automatically redirected to another warehouse. Multi-cluster scaling occurs within the same virtual warehouse.

Official Snowflake documentation reference:

Snowflake documentation explains that multi-cluster warehouses can automatically start and stop additional clusters to manage query concurrency. In Auto-scale mode, Snowflake starts additional clusters as needed, up to the maximum number configured for the warehouse.

Reference: Snowflake Documentation - Multi-cluster warehouses; Snowflake Documentation - Auto-scale mode; SnowPro Core Study Guide - Virtual Warehouses and Resource Management.

NEW QUESTION: 30

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: ([SHOW ANSWER](#))

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: 31

Which statement describes pruning?

A. The filtering or disregarding of micro-partitions that are not needed to return a query.

B. The return of micro-partitions values that overlap with each other to reduce a query 's runtime.

C. A service that is handled by the Snowflake Cloud Services layer to optimize caching.

D. The ability to allow the result of a query to be accessed as if it were a table.

Answer: ([**SHOW ANSWER**](#)**)**

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

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here: <https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (**1187** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 32

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: ([**SHOW ANSWER**](#)**)**

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 UDF.

NEW QUESTION: 33

Which types of values lead to slower queries and increased storage consumption when placed within a VARIANT column? Select TWO.

- A. JSON null values
- B. Floating-point numbers
- C. Arrays
- D. Vectors
- E. Timestamps within strings

Answer: ([**SHOW ANSWER**](#)**)**

The correct answers are A. JSON null values and E. Timestamps within strings .

Important correction:

The provided answer A, B is not the best answer. Snowflake documentation specifically calls out that certain values in semi-structured data may not be extracted into separate columnar storage, which can result in slower queries and increased storage usage. Two commonly tested examples are JSON null values and date or timestamp values stored as strings.

Why A is correct:

When a semi-structured element contains JSON null values, Snowflake might not extract that element into a separate column. This can reduce pruning efficiency and require Snowflake to scan more of the VARIANT structure during query execution.

Why E is correct:

Date and timestamp values inside semi-structured data that are stored as strings may not be optimized in the same way as native typed values. Storing timestamps as strings can increase storage and slow query performance because Snowflake may need to parse or cast the values at query time.

Why the other options are incorrect:

B). Floating-point numbers are supported inside VARIANT; they are not the main documented issue in this context.

C). Arrays are supported in semi-structured data. While very complex nested structures can affect query patterns, arrays alone are not the best answer here.

D). Vectors are a Snowflake data type used for vector operations and are not the standard answer for this semi-structured data optimization question.

Official Snowflake documentation reference:

Snowflake documentation explains that for semi-structured data, elements containing JSON null values may not be extracted into separate columns. It also recommends storing date and timestamp values in native typed columns when possible rather than as strings inside VARIANT, to improve query performance and storage efficiency.

Reference: Snowflake Documentation - Considerations for semi-structured data stored in VARIANT; Snowflake Documentation - Semi-structured data types; SnowPro Core Study Guide - Working with Semi- Structured Data.

NEW QUESTION: 34

What is the expiration period for a file URL used to access unstructured data in cloud storage?

- A. An unlimited amount of time
- B. The remainder of the session
- C. The same length of time as the expiration period for the query results cache
- D. The length of time specified in the expiration_time argument

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 35

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 36

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: ([SHOW ANSWER](#))

To show all privileges granted to a specific schema, the command SHOW GRANTS ON SCHEMA < schema_name > should be used. In this case, it would be SHOW GRANTS ON SCHEMA ANALYTICS_DW.MARKETING. References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 37

What privileges are required to allow role RL_CUST_WH to start warehouse CUST_WH?

- A. GRANT OPERATE ON CUST_WH TO ROLE RL_CUST_WH;
- B. GRANT CREATE WAREHOUSES ON ACCOUNT TO ROLE RL_CUST_WH;
- C. GRANT MANAGE WAREHOUSES ON ACCOUNT TO ROLE RL_CUST_WH;
- D. GRANT USAGE ON CUST_WH TO ROLE RL_CUST_WH;

Answer: ([SHOW ANSWER](#))

To start, resume, or suspend a warehouse, the role needs the OPERATE privilege. USAGE only allows reference to the object but not its activation.

Required Privileges for Starting a Warehouse:

- * USAGE (to reference)
- * OPERATE (to start/stop)

Reference:

Snowflake Docs: Warehouse Privileges

NEW QUESTION: 38

A department uses a multi-cluster virtual warehouse to run ad hoc reports for multiple business analysts. The warehouse load history shows that many report queries were queuing before they were run. How can this issue be solved?

- A. Set up a schedule to limit how many reports can be run at the same time.
- B. Add more clusters to the warehouse.
- C. Increase the size of the warehouse.
- D. Convert the warehouse to a Snowpark-optimized virtual warehouse.

Answer: B ([LEAVE A REPLY](#))

The correct answer is B. Add more clusters to the warehouse .

Queued queries on a multi-cluster warehouse usually indicate a concurrency problem. Multiple analysts are submitting queries at the same time, and the warehouse does not have enough active cluster capacity to process them concurrently.

Why B is correct:

Adding more clusters to a multi-cluster warehouse increases concurrency capacity. In Auto-scale mode, Snowflake can start additional clusters as needed, up to the configured maximum, to reduce queuing.

Why the other options are incorrect:

A). Scheduling or limiting reports may reduce concurrency, but it adds operational overhead and does not use Snowflake's built-in multi-cluster scaling capability.

C). Increasing warehouse size can make individual queries faster, but it does not directly improve concurrency as effectively as adding clusters.

D). A Snowpark-optimized warehouse is intended for memory-intensive Snowpark workloads, not for reducing query queues from many concurrent BI/ad hoc users.

Official Snowflake documentation reference:

Snowflake documentation explains that multi-cluster warehouses are designed to handle concurrency by automatically scaling out with additional clusters.

Reference: Snowflake Documentation - Multi-cluster warehouses; Warehouse load monitoring; SnowPro Core Study Guide - Snowflake Account and Resource Management.

NEW QUESTION: 39

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: ([SHOW ANSWER](#))

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

NEW QUESTION: 40

A virtual warehouse's auto-suspend and auto-resume settings apply to which of the following?

A. The primary cluster in the virtual warehouse

B. The entire virtual warehouse

C. The database in which the virtual warehouse resides

D. The Queries currently being run on the virtual warehouse

Answer: ([SHOW ANSWER](#))

The auto-suspend and auto-resume settings in Snowflake apply to the entire virtual warehouse. These settings allow the warehouse to automatically suspend when it's not in use, helping to save on compute costs. When queries or tasks are submitted to the warehouse, it can automatically resume operation. This functionality is designed to optimize resource usage and cost-efficiency.

References:

SnowPro Core Certification Exam Study Guide (as of 2021)

Snowflake documentation on virtual warehouses and their settings (as of 2021)

NEW QUESTION: 41

Based on a review of a Query Profile, which scenarios will benefit the MOST from the use of a data clustering key?
(Select TWO.)

- A. A column that appears most frequently in aggregate operations
- B. A column that appears most frequently in where operations
- C. A column that appears most frequently in join operations
- D. A column that appears most frequently in order by operations
- E. A column that appears most frequently in group by operations

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 42

Which file format option should be used when unloading data into a stage to create a CSV or a JSON file?

- A. FILE_EXTENSION
- B. PARSE_HEADER
- C. TRIM_SPACE
- D. SKIP_HEADER

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 43

What action will cause the cache associated with compute resources to be dropped?

- A. Decreasing the size of a running standard virtual warehouse
- B. Adding a multi-cluster virtual warehouse that is of a similar size to a running warehouse
- C. Increasing the size of a running standard virtual warehouse
- D. Adding a new standard virtual warehouse that is the same size as a running warehouse

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 44

Which is the MINIMUM required Snowflake edition that a user must have if they want to use AWS/Azure Privatelink or Google Cloud Private Service Connect?

- A. Standard
- B. Premium
- C. Enterprise
- D. Business Critical

Answer: D ([LEAVE A REPLY](#))

<https://docs.snowflake.com/en/user-guide/admin-security-privatelink.html>

NEW QUESTION: 45

Which security feature supports the replacement of active keys to enhance key-pair authentication?

- A. Authentication
- B. Passkeys
- C. Rotation
- D. Time-based one-time passcodes

Answer: (SHOW ANSWER)

The correct answer is C. Rotation .

Key-pair authentication uses a public/private key pair to authenticate users or applications. To improve security, keys should be replaced periodically or when there is a possibility that a key has been exposed. This process is called key rotation .

Why C is correct:

Key rotation allows an organization to replace active keys with new keys. In Snowflake, users can have public keys assigned for key-pair authentication, and rotating keys helps maintain secure authentication practices.

Why the other options are incorrect:

A). Authentication is the overall process of verifying identity, not the specific feature for replacing keys.

B). Passkeys are a different authentication mechanism and are not the feature used to replace active key-pair authentication keys.

D). Time-based one-time passcodes are associated with MFA, not key-pair replacement.

Official Snowflake documentation reference:

Snowflake documentation describes key-pair authentication and includes key rotation as a security practice for replacing public/private key pairs.

Reference: Snowflake Documentation - Key-pair authentication; Snowflake Documentation - Key rotation; SnowPro Core Study Guide - Security and Access Control.

NEW QUESTION: 46

Which of the following is a data tokenization integration partner?

A. Protegrity

B. Tableau

C. DBeaver

D. SAP

Answer: (SHOW ANSWER)

Protegrity is listed as a data tokenization integration partner for Snowflake. This partnership allows Snowflake users to utilize Protegrity's tokenization solutions within the Snowflake environment³.

References = [COF-C03] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 47

Using which object-level parameters will help limit query processing and concurrency slowdowns? (Select TWO).

A. STATEMENT_TIMEOUT_IN_SECONDS

- B. STATEMENT_QUEUED_TIMEOUT_IN_SECONDS
- C. CHANGE_TRACKING
- D. DATA_RETENTION_TIME_IN_DAYS
- E. MAX_DATA_EXTENSION_TIME_IN_DAYS
- F. QUERY_TAG

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 48

Which Snowflake objects can be restored using Time Travel? (Select VNO).

- A. Roles
- B. Users
- C. Databases
- D. Schemas
- E. Virtual warehouses

Answer: ([SHOW ANSWER](#))

Snowflake's Time Travel feature allows users to access historical data within a specific period. This feature supports the restoration of various objects, including databases and schemas, to their previous states. Time Travel can be used for recovering dropped objects, undoing accidental changes, or analyzing data changes over time. However, it does not support user or role objects like Users and Roles, or compute resources like Virtual Warehouses.

References: Snowflake Documentation on Time Travel

NEW QUESTION: 49

When connecting to Snowflake using SnowSQL, what are ways to explicitly specify the password? (Select TWO).

- A. Enter through an interactive prompt
- B. Use public and private key pair authentication
- C. Use an OAuth token
- D. Specify using SNOWSQL_PWD environment variables
- E. Run a web-based authorization flow

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 50

Based on Snowflake recommendations, when creating a hierarchy of custom roles, the top-most custom role should be assigned to which role?

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

Answer: ([SHOW ANSWER](#))

Based on Snowflake recommendations, when creating a hierarchy of custom roles, the top-most custom role should ideally be granted to the ACCOUNTADMIN role. This recommendation stems from the best practices for implementing a least privilege access control model, ensuring that only the necessary permissions are granted at each level of the

role hierarchy. The ACCOUNTADMIN role has the highest level of privileges in Snowflake, including the ability to manage all aspects of the Snowflake account. By assigning the top-most custom role to ACCOUNTADMIN, you ensure that the administration of role hierarchies and the assignment of roles remain under the control of users with the highest level of oversight and responsibility within the Snowflake environment.

References:

Snowflake Documentation on Access Control: Managing Access Control

NEW QUESTION: 51

When a schema is cloned what happens to the tags that reside in that schema or database?

- A. The tags remain in the source schema unchanged.
- B. The tags are deleted and must be redefined in the target schema.
- C. The tags are automatically cloned with the schema, maintaining all associations.
- D. The tags become temporarily inactive until they are manually reactivated.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 52

Which setting will help control credit consumption for a virtual warehouse that needs to run in Auto-scale mode?

- A. Set the warehouse scaling policy to ECONOMY.
- B. Set the MIN_CLUSTER_COUNT to 1.
- C. Set the AUTO_SUSPEND property of the warehouse to 0.
- D. Set the MIN_CLUSTER_COUNT to 2 and set the MAX_CLUSTER_COUNT to 2.

Answer: ([SHOW ANSWER](#))

The correct answer is A. Set the warehouse scaling policy to ECONOMY .

A multi-cluster warehouse can use Auto-scale mode to add or remove clusters based on query concurrency.

To help control credit consumption, Snowflake provides an Economy scaling policy.

Why A is correct:

The ECONOMY scaling policy is more conservative than the STANDARD scaling policy. It starts additional clusters only when Snowflake estimates there is enough query load to keep the new cluster busy for a meaningful amount of time. This helps reduce unnecessary cluster starts and can lower credit usage.

Why the other options are incorrect:

- B). Setting MIN_CLUSTER_COUNT = 1 can help avoid running multiple clusters constantly, but it does not specifically control Auto-scale behavior as effectively as the Economy scaling policy.
- C). Setting AUTO_SUSPEND = 0 disables automatic suspension, which can increase credit consumption.
- D). Setting both minimum and maximum cluster count to 2 forces two clusters and removes Auto-scale flexibility. This can increase cost and does not control consumption efficiently.

Official Snowflake documentation reference:

Snowflake documentation explains that the Economy scaling policy conserves credits by starting additional clusters only when enough query load exists to justify them.

Reference: Snowflake Documentation - Multi-cluster warehouses; Snowflake Documentation - Scaling policy; SnowPro Core Study Guide - Snowflake Account and Resource Management.

NEW QUESTION: 53

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

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

Answer: ([SHOW ANSWER](#))

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

References:

Snowflake Documentation: Data Replication

NEW QUESTION: 54

Which native data types are used for storing semi-structured data in Snowflake? (Select TWO)

- A. NUMBER
- B. OBJECT
- C. STRING
- D. VARCHAR
- E. VARIANT

Answer: ([SHOW ANSWER](#))

Snowflake supports semi-structured data types, which include OBJECT and VARIANT. These data types are capable of storing JSON-like data structures, allowing for flexibility in data representation. OBJECT can directly contain VARIANT, and thus indirectly contain any other data type, including itself¹.

NEW QUESTION: 55

In which hierarchy is tag inheritance possible?

- A. Organization » Account » Role
- B. Account » User » Schema
- C. Database » View » Column
- D. Schema » Table » Column

Answer: ([SHOW ANSWER](#))

In Snowflake, tag inheritance is a feature that allows tags, which are key-value pairs assigned to objects for the purpose of data governance and metadata management, to be inherited within a hierarchy. The hierarchy in which tag inheritance is possible is from Schema to Table to Column. This means that a tag applied to a schema can be inherited by the tables within that schema, and a tag applied to a table can be inherited by the columns within that table.

References: Snowflake Documentation on Tagging and Object Hierarchy

NEW QUESTION: 56

The effects of query pruning can be observed by evaluating which statistics? (Select TWO).

- A. Partitions scanned
- B. Partitions total
- C. Bytes scanned
- D. Bytes read from result
- E. Bytes written

Answer: ([SHOW ANSWER](#))

Query pruning in Snowflake refers to the optimization technique where the system reduces the amount of data scanned by a query based on the query conditions. This typically involves skipping unnecessary data partitions that do not contribute to the query result. The effectiveness of this technique can be observed through:

Option A: Partitions scanned. This statistic indicates how many data partitions were actually scanned as a result of query pruning, showing the optimization in action.

Option C: Bytes scanned. This measures the volume of data physically read during query execution, and a reduction in this number indicates effective query pruning, as fewer bytes are read when unnecessary partitions are skipped.

Options B, D, and E do not directly relate to observing the effects of query pruning. "Partitions total" shows the total available, not the impact of pruning, while "Bytes read from result" and "Bytes written" relate to output rather than the efficiency of data scanning.

References: Snowflake documentation on performance tuning and query optimization techniques, specifically how query pruning affects data access.

NEW QUESTION: 57

Which Snowflake feature can be used to find sensitive data in a table or column?

- A. Masking policies
- B. Data classification
- C. Row level policies
- D. External functions

Answer: ([SHOW ANSWER](#))

Data classification in Snowflake is a feature that allows organizations to identify and categorize data stored in tables or columns based on its sensitivity level or content type. This feature can be used to find sensitive data within the database by classifying data as confidential, personal, public, etc., making it easier to apply appropriate security measures, such as masking policies or row-level security, to protect sensitive information.

References:

Snowflake Documentation: Data Classification

NEW QUESTION: 58

Which Snowflake table is an implicit object layered on a stage, where the stage can be either internal or external?

- A. Directory table
- B. Temporary table
- C. Transient table

D. A table with a materialized view

Answer: ([SHOW ANSWER](#))

A directory table in Snowflake is an implicit object layered on a stage, whether internal or external. It allows users to query the contents of a stage as if it were a table, providing metadata about the files stored in the stage, such as filenames, file sizes, and last modified timestamps.

Snowflake Documentation: Directory Tables

NEW QUESTION: 59

A user with which privileges can create or manage other users in a Snowflake account? (Select TWO).

A. GRANT

B. SELECT

C. MODIFY

D. OWNERSHIP

E. CREATE USER

Answer: ([SHOW ANSWER](#))

A user with the OWNERSHIP privilege on a user object or the CREATE USER privilege on the account can create or manage other users in a Snowflake account56.

NEW QUESTION: 60

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: ([SHOW ANSWER](#))

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: 61

Which option will create a table without data from the CUSTOMERS table?

A. CREATE TABLE CUSTOMERS_CLONE LIKE CUSTOMERS;

B. CREATE TABLE CUSTOMERS_CLONE CLONE CUSTOMERS;

C. CREATE OR REPLACE TEMP TABLE CUSTOMERS_CLONE (CLONE VARCHAR);

D. CREATE TABLE CUSTOMERS_CLONE AS SELECT * FROM CUSTOMERS;

Answer: ([SHOW ANSWER](#))

The correct answer is A. CREATE TABLE CUSTOMERS_CLONE LIKE CUSTOMERS; .

The CREATE TABLE ... LIKE command creates a new table with the same column definitions as an existing table, but it does not copy the data from the source table.

Why A is correct:

CREATE TABLE CUSTOMERS_CLONE LIKE CUSTOMERS; creates an empty table with the same structure as CUSTOMERS.

Why the other options are incorrect:

B). CREATE TABLE ... CLONE creates a zero-copy clone of the source table, including access to the source table's existing data at the time of cloning.

C). This creates a temporary table with one column named CLONE; it does not create a table based on the structure of CUSTOMERS.

D). CREATE TABLE ... AS SELECT creates a new table and inserts the result of the query, so it copies data from CUSTOMERS.

Official Snowflake documentation reference:

Snowflake documentation explains that CREATE TABLE ... LIKE creates a new table with the same column definitions as another table, but without copying data.

Reference: Snowflake Documentation - CREATE TABLE ... LIKE; Snowflake Documentation - CREATE TABLE ... CLONE; SnowPro Core Study Guide - SQL and Snowflake Objects.

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:

<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF Special Discount**

Code: **freecram**)

NEW QUESTION: 62

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: (SHOW ANSWER)

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: 63

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: ([SHOW ANSWER](#))

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: 64

Which privilege is required on a virtual warehouse to abort any existing executing queries?

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

Answer: ([SHOW ANSWER](#))

The privilege required on a virtual warehouse to abort any existing executing queries is OPERATE.

The OPERATE privilege on a virtual warehouse allows a user to perform operational tasks on the warehouse, including starting, stopping, and restarting the warehouse, as well as aborting running queries. This level of control is essential for managing resource utilization and ensuring that the virtual warehouse operates efficiently.

References:

Snowflake Documentation on Access Control: Access Control Privileges

NEW QUESTION: 65

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 ([LEAVE A REPLY](#))

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: 66

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: ([SHOW ANSWER](#))

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: 67

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 68

Which metrics in the QUERY_HISTORY Account_Usage View can be used to assess the pruning efficiency of a query? (Select TWO).

- A. COMPILATION_TIME
- B. PARTITIONS_TOTAL
- C. EXECUTION_TIME
- D. TOTAL_ELAPSED_TIME
- E. PARTITIONS_SCANNED

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 69

Who can activate a network policy for users in a Snowflake account? (Select TWO)

- A. ACCOUNTADMIN
- B. USERADMIN
- C. PUBLIC
- D. SYSADMIN
- E. Any role that has the global ATTACH POLICY privilege

Answer: ([SHOW ANSWER](#))

Network policies in Snowflake are used to control access to Snowflake accounts based on IP address ranges. These policies can be activated by specific roles that have the necessary privileges.

Role: ACCOUNTADMIN:

The ACCOUNTADMIN role has full administrative rights across the Snowflake account.
This role can manage all aspects of the Snowflake environment, including network policies.
Role with Global ATTACH POLICY Privilege:
Any role that has been granted the global ATTACH POLICY privilege can activate network policies.
This privilege allows the role to attach policies that control network access to the account.
Snowflake Documentation: Network Policies

NEW QUESTION: 70

When used with a SELECT statement in the COPY INTO <location> command, what statement will unload data from a table as JSON?

- A. PARSE_JSON
- B. OBJECT_CONSTRUCT
- C. TO_OBJECT
- D. TO_JSON

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 71

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

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

Answer: ([SHOW ANSWER](#))

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

References:

Snowflake Documentation on JSON File Format

[COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 72

When used with the UNLOAD command, which parameter specifies the destination of unloaded data?

- A. COPY INTO <table name>
- B. COPY INTO <stage name>
- C. GET <file name>
- D. PUT <file name>

Answer: ([SHOW ANSWER](#))

In Snowflake, the COPY INTO <stage name> syntax is used with the UNLOAD command to specify the target location where the data should be unloaded, typically a stage or cloud storage (such as Amazon S3 or Azure Blob Storage). This command unloads data from a Snowflake table into files within the specified destination, enabling easy export and external storage of data. GET and PUT commands are used for file management but are not related to unloading table data directly.

NEW QUESTION: 73

What actions can be taken to reduce memory spilling for queries that are too large to fit into memory? Select TWO.

- A.** Increase the size of the virtual warehouse.
- B.** Edit the queries to process the data in smaller batches.
- C.** Increase the size of the virtual warehouse clusters.
- D.** Suspend and then resume the virtual warehouse.
- E.** Set a clustering key on the tables used in the queries.

Answer: ([SHOW ANSWER](#))

The correct answers are A. Increase the size of the virtual warehouse and B. Edit the queries to process the data in smaller batches .

Memory spilling occurs when a query requires more memory than is available in the virtual warehouse.

Snowflake may spill intermediate data to local or remote storage, which can slow query performance.

Why A is correct:

Increasing the warehouse size provides more compute and memory resources. This can reduce or eliminate spilling for memory-intensive queries.

Why B is correct:

Rewriting queries to process data in smaller batches reduces the amount of intermediate data held in memory at one time. This can reduce spilling and improve performance.

Important note about option C:

Option C is misleading because a virtual warehouse cluster's size is changed by increasing the warehouse size . Since option A already says this correctly, option C is not the best answer.

Why the other options are incorrect:

C). "Increase the size of the virtual warehouse clusters" is essentially a less precise duplicate of increasing warehouse size. The exam-style answer is A.

D). Suspending and resuming a warehouse may clear cache, but it does not increase available memory and does not solve memory spilling.

E). Clustering can improve pruning and reduce scanned data for some queries, but it is not the direct general solution for queries that are spilling because they are too large for memory.

Official Snowflake documentation reference:

Snowflake Query Profile documentation explains that spilling to local or remote storage can occur when operations do not fit in memory. Snowflake performance guidance recommends scaling up the warehouse or rewriting queries to reduce the amount of data processed.

Reference: Snowflake Documentation - Query Profile; Snowflake Documentation - Warehouse performance; SnowPro Core Study Guide - Performance Optimization.

NEW QUESTION: 74

Which COPY INTO < table > parameter is used to avoid loading a partial dataset?

- A. FORCE = FALSE
- B. RETURN_FAILED_ONLY = FALSE
- C. LOAD_UNCERTAIN_FILES = FALSE
- D. ON_ERROR = ABORT_STATEMENT

Answer: (SHOW ANSWER)

The correct answer is D. ON_ERROR = ABORT_STATEMENT .

ON_ERROR = ABORT_STATEMENT stops the load operation if an error is encountered. This helps avoid loading only part of a dataset when a file or row error occurs.

Why D is correct:

If the goal is to avoid a partial dataset, the load should fail rather than continue after errors.

ABORT_STATEMENT prevents Snowflake from loading only the valid rows while skipping invalid ones.

Why the other options are incorrect:

- A). FORCE = FALSE prevents reloading files that were already loaded, but it does not prevent partial loads caused by data errors.
- B). RETURN_FAILED_ONLY affects command output reporting. It does not control whether partial data is loaded.
- C). LOAD_UNCERTAIN_FILES = FALSE controls loading files with uncertain load status, but it is not the primary parameter for avoiding partial datasets caused by load errors.
- D). ON_ERROR = ABORT_STATEMENT is the correct parameter because it stops the load when an error occurs.

Official Snowflake documentation reference:

Snowflake documentation for COPY INTO < table > describes ON_ERROR = ABORT_STATEMENT as the option that aborts the load operation when an error is encountered.

Reference: Snowflake Documentation - COPY INTO < table > ; Snowflake Documentation - ON_ERROR copy option; SnowPro Core Study Guide - Data Loading and Unloading.

NEW QUESTION: 75

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: (SHOW ANSWER)

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

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

- A. Masking policy

- B. Stored procedure
- C. Task
- D. External table

Answer: (SHOW ANSWER)

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

References:

Snowflake Documentation on Secure Data Sharing

SnowPro Core Certification Companion: Hands-on Preparation and Practice

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (**1187** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 77

What is the purpose of setting the Query Acceleration Service scale factor?

- A. To establish the maximum number of clustering keys that can be created by the Query Acceleration Service
- B. To limit the number of users who can use the Query Acceleration Service
- C. To establish the maximum virtual warehouse size that can be used by the Query Acceleration Service
- D. To control the potential costs associated with the Query Acceleration Service

Answer: D (LEAVE A REPLY)

The correct answer is D. To control the potential costs associated with the Query Acceleration Service .

The Query Acceleration Service can use additional Snowflake-managed compute resources to accelerate eligible queries. The scale factor setting controls the maximum amount of additional compute that can be used relative to the warehouse.

Why D is correct:

A higher scale factor allows more acceleration resources to be used, which can improve performance but may increase credit consumption. A lower scale factor limits the amount of extra compute used, helping control cost.

Example concept:

```
ALTER WAREHOUSE my_wh
SET ENABLE_QUERY_ACCELERATION = TRUE
QUERY_ACCELERATION_MAX_SCALE_FACTOR = 4;
```

This limits the maximum acceleration scale factor for the warehouse.

Why the other options are incorrect:

A). Query Acceleration Service does not create clustering keys.

- B). The scale factor does not control how many users can use the service.
- C). The scale factor does not resize the virtual warehouse itself. It limits additional query acceleration resources.

Official Snowflake documentation reference:

Snowflake documentation explains that QUERY_ACCELERATION_MAX_SCALE_FACTOR limits the amount of compute resources that can be used by the Query Acceleration Service for a warehouse, which helps control cost.

Reference: Snowflake Documentation - Query Acceleration Service; Snowflake Documentation -

QUERY_ACCELERATION_MAX_SCALE_FACTOR; SnowPro Core Study Guide - Performance Optimization.

NEW QUESTION: 78

Which function, when combined with a COPY INTO < location > command, will convert the rows in a relational table to a single VARIANT column while the rows are being unloaded?

- A. ARRAY_AGG
- B. OBJECT_CONSTRUCT
- C. TO_VARIANT
- D. OBJECT_AGG

Answer: (SHOW ANSWER)

The correct answer is B. OBJECT_CONSTRUCT .

When unloading relational table data into semi-structured formats such as JSON, OBJECT_CONSTRUCT can be used to convert each relational row into an object. The resulting object can be represented as a single VARIANT value.

Why B is correct:

OBJECT_CONSTRUCT creates an object from key-value pairs. When used with *, it can construct an object from all columns in a row.

Example concept:

```
COPY INTO @my_stage/output/  
FROM (  
  SELECT OBJECT_CONSTRUCT(*)  
  FROM my_table  
)  
FILE_FORMAT = (TYPE = JSON);
```

This transforms each relational row into a JSON-style object during unload.

Why the other options are incorrect:

- A). ARRAY_AGG aggregates multiple values into an array, but it does not convert each relational row into an object.
- C). TO_VARIANT converts a value to VARIANT, but it does not automatically construct a structured object from all row columns.
- D). OBJECT_AGG is an aggregate function that creates an object from key-value pairs across rows, not the usual function for converting each relational row into a single object during unload.

Official Snowflake documentation reference:

Snowflake documentation explains that OBJECT_CONSTRUCT returns an object constructed from key- value pairs and can be used with COPY INTO < location > to unload relational rows as semi-structured data.

Reference: Snowflake Documentation - OBJECT_CONSTRUCT; Snowflake Documentation - Unloading relational data to JSON; Snowflake Documentation - COPY INTO < location > ; SnowPro Core Study Guide - Working with Semi-Structured Data.

NEW QUESTION: 79

Which semi-structured file format is a compressed, efficient, columnar data representation?

- A. Avro
- B. JSON
- C. TSV
- D. Parquet

Answer: ([SHOW ANSWER](#))

Parquet is a columnar storage file format that is optimized for efficiency in both storage and processing. It supports compression and encoding schemes that significantly reduce the storage space needed and speed up data retrieval operations, making it ideal for handling large volumes of data. Unlike JSON or TSV, which are row-oriented and typically uncompressed, Parquet is designed specifically for use with big data frameworks, offering advantages in terms of performance and cost when storing and querying semi-structured data.

References: Apache Parquet Documentation

NEW QUESTION: 80

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 81

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 82

What type of queries will benefit from the search optimization service when using equality predicates?

- A. Queries using data substrings
- B. Queries with geospatial data
- C. Queries using point-lookups
- D. Queries of semi-structured data

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 83

Which function will generate a URL that can be accessed by a non-Snowflake user?

- A. GET_PREIGNED_URL
- B. GET_STAGE_LOCATION
- C. BUILD_STAGE_FILE_URL
- D. BUILD_SCOPED_FILE_URL

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 84

Which step will increase security related to user access to a Snowflake account?

- A. Set up a User-Defined Function, or UDF, to run a daily report identifying all security-related session parameters used in the last 24 hours.
- B. Restrict access to Snowflake to known networks using a network policy.
- C. Establish a very short session timeout parameter.
- D. Default all new users to a PUBLIC role with restrictive permissions.

Answer: ([SHOW ANSWER](#))

The correct answer is B. Restrict access to Snowflake to known networks using a network policy .

A network policy improves account access security by allowing or blocking access based on IP address ranges. This helps ensure that users can connect to Snowflake only from approved networks.

Why B is correct:

Network policies can restrict inbound access to Snowflake by defining allowed and blocked IP address ranges. This is a direct security control for user access to a Snowflake account.

Example concept:

```
CREATE NETWORK POLICY corp_network_policy  
ALLOWED_IP_LIST = ( ' 203.0.113.0/24 ' );  
ALTER ACCOUNT SET NETWORK_POLICY = corp_network_policy;
```

Why the other options are incorrect:

A). A UDF is not the correct tool for managing account access security. Monitoring security parameters may provide visibility, but it does not directly restrict access.

C). A short session timeout can reduce risk from idle sessions, but making it "very short" can harm usability and is not as direct as network restrictions.

D). Every user already has access to the PUBLIC role by default. Setting restrictive permissions on PUBLIC is good practice, but it is not as direct as restricting access through a network policy.

Official Snowflake documentation reference:

Snowflake documentation describes network policies as security controls that restrict access to Snowflake based on allowed and blocked IP address lists.

Reference: Snowflake Documentation - Network policies; Snowflake Documentation - Network policy best practices; SnowPro Core Study Guide - Security and Access Control.

NEW QUESTION: 85

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

A. READ

B. OWNERSHIP

C. USAGK

D. WRTTF

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 86

What AUTHENTICATION_METHODS value should be selected for the Single Sign-On, or SSO, login option to appear when creating a Snowflake authentication policy?

A. PASSWORD

B. KEYPAIR

C. SAML

D. OAUTH

Answer: ([SHOW ANSWER](#))

The correct answer is C. SAML .

In Snowflake, Single Sign-On is commonly implemented using federated authentication through a SAML 2.0 identity provider. When creating an authentication policy, including SAML in the allowed AUTHENTICATION_METHODS enables SSO-based login options.

Why C is correct:

SAML represents SAML-based federated authentication. This is the authentication method associated with SSO login in Snowflake.

Example concept:

```
CREATE AUTHENTICATION POLICY sso_policy  
AUTHENTICATION_METHODS = ( ' SAML ' );
```

Why the other options are incorrect:

A). PASSWORD allows username and password authentication, not SSO.

B). KEYPAIR allows key-pair authentication, commonly used for service accounts and applications.

D). OAUTH allows OAuth-based authentication, but the SSO login option in this context is associated with SAML.

Official Snowflake documentation reference:

Snowflake documentation describes authentication policies as controlling allowed authentication methods, including password, SAML, OAuth, and key-pair authentication. SAML is used for federated SSO authentication.

Reference: Snowflake Documentation - Authentication policies; Snowflake Documentation - Federated authentication and SSO; SnowPro Core Study Guide - Security and Access Control.

NEW QUESTION: 87

Which Snowflake object can be used to record DML changes made to a table?

A. Snowpipe

B. Stage

C. Stream

D. Task

Answer: ([SHOW ANSWER](#))

Snowflake Streams are used to track and record Data Manipulation Language (DML) changes made to a table.

Streams capture changes such as inserts, updates, and deletes, which can then be processed by other Snowflake objects or external applications.

Creating a Stream:

```
CREATE OR REPLACE STREAM my_stream ON TABLE my_table;
```

Using Streams: Streams provide a way to process changes incrementally, making it easier to build efficient data pipelines.

Consuming Stream Data: The captured changes can be consumed using SQL queries or Snowflake tasks.

References:

Snowflake Documentation: Using Streams

Snowflake Documentation: Change Data Capture (CDC) with Streams

NEW QUESTION: 88

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: ([SHOW ANSWER](#))

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.

Snowflake Documentation: Network Policy Management

NEW QUESTION: 89

A Snowflake account holds data that will be consumed by a BI tool. Most of the BI users are located in the United States, but some users will require access from abroad with different levels of concurrency.

What virtual warehouse would be MOST cost-efficient?

- A. A large enough standard warehouse able to handle the expected maximum concurrency.
- B. A maximized multi-cluster warehouse with the minimum number of clusters that are big enough to handle the expected maximum concurrency.
- C. An auto-scale, multi-cluster warehouse with a maximum number of clusters that are large enough to handle the expected maximum concurrency.
- D. A large enough Snowpark-optimized warehouse that will have sufficient memory to handle the expected maximum concurrency.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 90

Which command will indicate whether a materialized view is being used in a query?

- A. EXPLAIN
- B. DESCRIBE
- C. SHOW
- D. USE

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 91

What is the Fail-safe period for a transient table in the Snowflake Enterprise edition and higher?

- A. 0 days
- B. 1 day
- C. 7 days
- D. 14 days

Answer: ([SHOW ANSWER](#))

The Fail-safe period for a transient table in Snowflake, regardless of the edition (including Enterprise edition and higher), is 0 days. Fail-safe is a data protection feature that provides additional retention beyond the Time Travel period for recovering data in case of accidental deletion or corruption. However, transient tables are designed for temporary or short-term use and do not benefit from the Fail-safe feature, meaning that once their Time Travel period expires, data cannot be recovered.

References:

Snowflake Documentation: Understanding Fail-safe

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:

Code: **freecram**)

NEW QUESTION: 92

How does a Snowflake user extract the URL of a directory table on an external stage for further transformation?

- A. Use the SHOW STAGES command.
- B. Use the DESCRIBE STAGE command.
- C. Use the GET_ABSOLUTE_PATH function.
- D. Use the GET_STAGE_LOCATION function.

Answer: C ([LEAVE A REPLY](#))

To extract the URL of a directory table on an external stage for further transformation in Snowflake, the GET_ABSOLUTE_PATH function can be used. This function returns the full path of a file or directory within a specified stage, enabling users to dynamically construct URLs for accessing or processing data stored in external stages.

References:

Snowflake Documentation: Working with Stages

NEW QUESTION: 93

A user needs to automatically retrieve Iceberg column definitions using the INFER_SCHEMA function for semi-structured files stored on a stage. Which file format configuration will help them achieve this?

- A. CREATE FILE FORMAT my_ff TYPE = ORC;
- B. CREATE FILE FORMAT my_ff TYPE = AVRO;
- C. CREATE FILE FORMAT my_ff TYPE = PARQUET USE_VECTORIZED_SCANNER = TRUE;
- D. CREATE FILE FORMAT my_ff TYPE = PARQUET USE_VECTORIZED_SCANNER = FALSE;

Answer: ([SHOW ANSWER](#)**)**

For Parquet files, using the vectorized scanner supports improved schema inference behavior, including use cases involving Iceberg-compatible column definitions.

Corrected typos:

CREATE FIZE FORMAT should be CREATE FILE FORMAT.

VSE_VECTORIZED_SCANNER should be USE_VECTORIZED_SCANNER.

Reference: Snowflake Documentation - INFER_SCHEMA; Parquet file format options; Iceberg table schema inference.

NEW QUESTION: 94

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

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

Answer: ([SHOW ANSWER](#)**)**

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

NEW QUESTION: 95

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: ([SHOW ANSWER](#))

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: 96

Which command line parameter value can be pre-specified as an environment variable in SnowSQL?

- A. VARIABLE
- B. HOST
- C. OPTION
- D. MFA-PASSCODE

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 97

Which feature is used to protect sensitive data at query run time?

- A. Object tagging
- B. Data encryption
- C. User-Defined Functions, or UDFs
- D. Dynamic Data Masking

Answer: ([SHOW ANSWER](#))

The correct answer is D. Dynamic Data Masking .

Dynamic Data Masking is used to protect sensitive column values at query run time. It allows Snowflake to return different values depending on the user's role or other policy logic.

Why D is correct:

A masking policy can be applied to a table or view column. When a query runs, Snowflake evaluates the masking policy and determines whether to return the original value, a partially masked value, or a fully masked value.

Example concept:

```
CREATE MASKING POLICY phone_mask AS (val STRING)
```

```
RETURNS STRING - >
```

```
CASE
```

```
WHEN CURRENT_ROLE() IN ( ' CUSTOMER_SUPPORT_FULL ' ) THEN val
```

```
ELSE CONCAT( ' XXX-XXX- ' , RIGHT(val, 4))
```

```
END;
```

This protects sensitive phone numbers at query time without changing the underlying stored data.

Why the other options are incorrect:

- A). Object tagging classifies or labels data for governance, but it does not mask values at query runtime by itself.
- B). Data encryption protects data at rest and in transit, but it does not dynamically change query results based on role.
- C). UDFs can implement custom logic, but Snowflake's built-in governance feature for runtime masking is Dynamic Data Masking.

Official Snowflake documentation reference:

Snowflake documentation describes Dynamic Data Masking as a column-level security feature that uses masking policies to selectively mask plain-text data in query results at runtime.

Reference: Snowflake Documentation - Dynamic Data Masking; Snowflake Documentation - Masking policies; SnowPro Core Study Guide - Data Protection and Governance.

NEW QUESTION: 98

What objects can be cloned within Snowflake? (Select TWO).

- A. Schemas
- B. Users
- C. External tables
- D. Internal named stages
- E. External named stages

Answer: (SHOW ANSWER)

In Snowflake, cloning is available for certain types of objects, allowing quick duplication without copying data:

Schemas: These can be cloned, enabling users to replicate entire schema structures, including tables and views, for development or testing.

Internal named stages: These stages, used to store data files within Snowflake, can also be cloned, preserving configurations for data loading.

Users and external objects (like external stages or tables) cannot be cloned due to their dependency on external data and configurations outside Snowflake.

NEW QUESTION: 99

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

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

Answer: (SHOW ANSWER)

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

References = [COF-C03] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

NEW QUESTION: 100

Which file formats are supported for unloading data from Snowflake? (Choose two.)

- A. Avro
- B. JSON

- C. ORC
- D. XML
- E. Delimited (CSV, TSV, etc.)

Answer: ([SHOW ANSWER](#))

Snowflake supports unloading data in JSON and delimited file formats such as CSV and TSV. These formats are commonly used for data interchange and are supported by Snowflake for unloading operations

NEW QUESTION: 101

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: ([SHOW ANSWER](#))

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: 102

When a replication schedule is set, how does Snowflake manage refreshes for failover groups?

- A. Multiple refreshes will run concurrently.
- B. Only one refresh will run at a time.
- C. Refreshes will run regardless of role privileges.
- D. If the current refresh was delayed, the next scheduled refresh will be canceled.

Answer: ([SHOW ANSWER](#))

The correct answer is B. Only one refresh will run at a time .

Failover groups can be refreshed on a schedule to replicate supported account objects and database objects to target accounts. Snowflake does not run overlapping refresh operations for the same failover group.

Why B is correct:

If a scheduled refresh is still running, Snowflake does not start another concurrent refresh for the same failover group. Refresh operations are managed so that only one refresh runs at a time.

Why the other options are incorrect:

- A). Multiple refreshes for the same failover group do not run concurrently.
- C). Refreshes still require the appropriate privileges and configured replication/failover group permissions.
- D). A delayed or long-running refresh does not mean the next refresh is automatically canceled in the way stated. The key behavior is that refreshes do not overlap.

Official Snowflake documentation reference:

Snowflake documentation for replication and failover groups explains that scheduled refreshes are managed by Snowflake and that refresh operations do not run concurrently for the same group.

Reference: Snowflake Documentation - Replication and failover groups; Snowflake Documentation - Failover group refresh; SnowPro Core Study Guide - Snowflake Account and Resource Management.

NEW QUESTION: 103

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: ([SHOW ANSWER](#))

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: 104

What is the purpose of multi-cluster virtual warehouses?

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

Answer: ([SHOW ANSWER](#))

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

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

NEW QUESTION: 105

How does Snowflake define its approach to Discretionary Access Control (DAC)?

- A. A defined level of access to an object
- B. An entity in which access can be granted

- C. Each object has an owner, who can in turn grant access to that object.
- D. Access privileges are assigned to roles, which are in turn assigned to users.

Answer: (SHOW ANSWER)

Snowflake implements Discretionary Access Control (DAC) by using a role-based access control model. In this model, access privileges are not directly assigned to individual objects or users but are encapsulated within roles. These roles are then assigned to users, effectively granting them the access privileges contained within the role. This approach allows for granular control over database access, making it easier to manage permissions in a scalable and flexible manner.

References: Snowflake Documentation on Access Control

NEW QUESTION: 106

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

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

Answer: (SHOW ANSWER)

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

References:

[COF-C03] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Query Profile¹

Snowpro Core Certification Exam Flashcards²

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:

<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF Special Discount**

Code: freecram)

NEW QUESTION: 107

What are valid values for the FIELD_OPTIONALLY_ENCLOSED_BY option in the copy into < location > command used during data unloading? (Select TWO).

- A. Single quote character (')
- B. ' NONE '
- C. ' NULL '
- D. NULL

E. NONE

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 108

While running a query on a virtual warehouse in auto-scale mode, additional clusters are started immediately if which setting is configured?

- A. MIN_CLUSTER_COUNT is decreased and new_min_clusters is less than running_clusters
- B. MAX_CLUSTER_COUNT is increased and new_max_clusters is greater than running_clusters
- C. MAX_CLUSTER_COUNT is decreased and new_max_clusters is less than running_clusters
- D. MIN_CLUSTER_COUNT is increased and new_min_clusters is greater than running_clusters

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 109

Which categories are included in the execution time summary in a Query Profile? (Select TWO).

- A. Pruning
- B. Spilling
- C. Initialization
- D. Local Disk I/O
- E. Percentage of data read from cache

Answer: ([SHOW ANSWER](#))

In the execution time summary of a Query Profile in Snowflake, the categories included provide insights into various aspects of query execution. " Pruning " refers to the process by which Snowflake reduces the amount of data scanned by eliminating partitions of data that are not relevant to the query, thus improving performance. " Initialization " represents the time taken for query planning and setup before actual execution begins. These metrics are crucial for understanding and optimizing query performance.

Reference: Snowflake Documentation on the Query Profile, which outlines the different metrics and categories included in the execution summary: <https://docs.snowflake.com/en/user-guide/ui-query-profile.html#execution-summary>

NEW QUESTION: 110

What does Snowflake recommend a user do if they need to connect to Snowflake with a tool or technology that is not listed in Snowflake partner ecosystem?

- A. Use Snowflake ' s native API.
- B. Use a custom-built connector.
- C. Contact Snowflake Support for a new driver.
- D. Connect through Snowflake's JDBC or ODBC drivers

Answer: ([SHOW ANSWER](#))

If a user needs to connect to Snowflake with a tool or technology that is not listed in Snowflake ' s partner ecosystem, Snowflake recommends using its JDBC or ODBC drivers. These drivers provide a standard method of connecting from various tools and programming languages to Snowflake, offering wide compatibility and flexibility. By using these

drivers, users can establish connections to Snowflake from their applications, ensuring they can leverage the capabilities of Snowflake regardless of the specific tools or technologies they are using.

References: Snowflake Documentation on Client Drivers

NEW QUESTION: 111

Which table is designed to store a catalog of stage files in cloud storage?

- A. Dynamic table
- B. Hybrid table
- C. Iceberg table
- D. Directory table

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 112

Where can file formats be specified when unloading data from Snowflake? Select TWO.

- A. In a SELECT statement
- B. In a CREATE STAGE statement
- C. In a CREATE DATABASE statement
- D. In a CREATE INTEGRATION statement
- E. In a COPY INTO < location > statement

Answer: ([SHOW ANSWER](#))

The correct answers are B. In a CREATE STAGE statement and E. In a COPY INTO < location > statement .

When unloading data from Snowflake, file format options can be specified directly in the COPY INTO < location > command or associated with a named stage. Defining the file format in a stage is useful when repeated loads or unloads use the same format settings.

Why B is correct:

A stage can include a file format definition or reference. When data is unloaded to that stage, the file format associated with the stage can be used.

Example:

```
CREATE STAGE my_stage
FILE_FORMAT = (TYPE = CSV FIELD_DELIMITER = ' , ');
COPY INTO @my_stage
FROM my_table;
```

Why E is correct:

The COPY INTO < location > command can specify file format options directly.

Example:

```
COPY INTO @my_stage/output/
FROM my_table
FILE_FORMAT = (TYPE = CSV COMPRESSION = GZIP);
```

Why the other options are incorrect:

A). A SELECT statement defines the query result to unload, but file format options are not specified as part of a normal SELECT.

C). A database definition does not define unload file format options.

D). An integration defines access to external services or storage, not the file format used for unloading.

Official Snowflake documentation reference:

Snowflake documentation explains that file format options can be specified in the COPY INTO < location > command or as part of a named stage definition.

Reference: Snowflake Documentation - COPY INTO < location > ; Snowflake Documentation - CREATE STAGE;

Snowflake Documentation - File formats; SnowPro Core Study Guide - Data Loading and Unloading.

NEW QUESTION: 113

Masking policies are created at what level in Snowflake?

A. Table

B. Column

C. Schema

D. Database

Answer: ([SHOW ANSWER](#))

Masking policies in Snowflake are created and applied at the column level. These policies are used to obfuscate sensitive data by masking the values in a specific column. Masking policies can enforce data privacy and protection measures by dynamically masking the data based on the role of the user querying the data.

Snowflake Documentation: Dynamic Data Masking

NEW QUESTION: 114

Which command is used to upload data files from a local directory or folder on a client machine to an internal stage, for a specified table?

A. GET

B. PUT

C. CREATE STREAM

D. COPY INTO < location >

Answer: ([SHOW ANSWER](#))

To upload data files from a local directory or folder on a client machine to an internal stage in Snowflake, the PUT command is used. The PUT command takes files from the local file system and uploads them to an internal Snowflake stage (or a specified stage) for the purpose of preparing the data to be loaded into Snowflake tables.

Syntax Example:

PUT file:// < local_file_path > @ < stage_name > ;

This command is crucial for data ingestion workflows in Snowflake, especially when preparing to load data using the COPY INTO command.

Reference: Snowflake Documentation on Using the PUT Command: <https://docs.snowflake.com/en/sql-reference/sql/put.html>

NEW QUESTION: 115

How can a Data Exchange Administrator provide a user with account access to a Data Exchange?

A. Enable the IMPORT SHARE privilege and grant this privilege to the user.

- B. Grant the user the USERADMIN role.
- C. Create a new database for the Data Exchange and provide access to the user.
- D. Add the user to the Data Exchange.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 116

How should a data provider securely share Snowflake objects with a data consumer who does not have a Snowflake account?

- A. Create a reader account for the consumer.
- B. Give the consumer owner ' s rights on the provider ' s Snowflake account.
- C. Create and replicate a share, then give the consumer access to the replication.
- D. Unload the data into the consumer ' s cloud storage.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 117

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 118

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

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

Answer: ([SHOW ANSWER](#))

When deciding on a clustering key for a table, Snowflake recommends using the columns that are most actively used in the select filters. This is because clustering by these columns can improve the performance of queries that filter on these values, leading to more efficient scans and better overall query performance². References: [COF-C03]
SnowPro Core Certification Exam Study Guide

NEW QUESTION: 119

How are micro-partitions enabled on Snowflake tables?

- A. Micro-partitioning requires a cluster key on a table.
- B. Micro-partitioning requires the use of the search optimization service.

- C. Micro-partitioning is defined by the user when a table is created.
- D. Micro-partitioning is automatically performed on a table.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 120

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

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

Answer: ([SHOW ANSWER](#))

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

References = DROP STAGE, REMOVE

NEW QUESTION: 121

What is a responsibility of Snowflake's virtual warehouses?

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

Answer: ([SHOW ANSWER](#))

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

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (**1187** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 122

What Snowflake tool can be used to share data and metadata to one or more accounts in Snowflake?

- A. Direct shares
- B. Listings
- C. Secure views

D. Data Exchange

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 123

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: ([SHOW ANSWER](#))

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: 124

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: ([SHOW ANSWER](#))

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: 125

What does the Activity area of Snowsight allow users to do? (Select TWO).

- A. Schedule automated data backups.
- B. Explore each step of an executed query.
- C. Monitor queries executed by users in an account.
- D. Create and manage user roles and permissions.
- E. Access Snowflake Marketplace to find and integrate datasets.

Answer: ([SHOW ANSWER](#))

The Activity area of Snowsight, Snowflake's web interface, allows users to perform several important tasks related to query management and performance analysis. Among the options provided, the correct ones are:

B . Explore each step of an executed query: Snowsight provides detailed insights into query execution, including the ability to explore the execution plan of a query. This helps users understand how a query was processed, identify performance bottlenecks, and optimize query performance.

C . Monitor queries executed by users in an account: The Activity area enables users to monitor the queries that have been executed by users within the Snowflake account. This includes viewing the history of queries, their execution times, resources consumed, and other relevant metrics.

These features are crucial for effective query performance tuning and ensuring efficient use of Snowflake's resources.

References:

Snowflake Documentation on Snowsight: Using Snowsight

NEW QUESTION: 126

How can a 5 GB table be downloaded into a single file MOST efficiently?

- A. Keep the default MAX_FILE_SIZE to 16 MB
- B. Set the default MAX_FILE_SIZE to 5 GB.
- C. Set the SINGLE parameter to TRUE.
- D. Use a regular expression in the stage specifications of the COPY command.

Answer: ([SHOW ANSWER](#))

To download a 5 GB table into a single file most efficiently in Snowflake, you should set the SINGLE parameter to TRUE. This parameter ensures that the COPY INTO command outputs the result into a single file, regardless of the file size. This approach is more efficient than relying on the default MAX_FILE_SIZE setting, which would split the output into multiple files.

Snowflake Documentation: COPY INTO < location >

NEW QUESTION: 127

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

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

Answer: ([SHOW ANSWER](#))

The user property 'MINS TO BYPASS MFA' in Snowflake allows temporary bypass of MFA for a user, which can be set by an account administrator without contacting Snowflake Support2.

NEW QUESTION: 128

Who can activate and enforce a network policy for all users in a Snowflake account? (Select TWO).

- A. A user with an USERADMIN or higher role
- B. A user with a SECURITYADMIN or higher role
- C. A role that has been granted the ATTACH POLICY privilege
- D. A role that has the NETWORK_POLICY account parameter set
- E. A role that has the ownership of the network policy

Answer: ([SHOW ANSWER](#))

In Snowflake, a user with the SECURITYADMIN role or higher can activate and enforce a network policy for all users in an account. Additionally, a role that has ownership of the network policy can also activate and enforce it

NEW QUESTION: 129

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

- A. Clustering
- B. Indexing
- C. Fragmenting
- D. Using INDEX__HINTS

Answer: A ([LEAVE A REPLY](#))

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

References:

[COF-C03] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Clustering

NEW QUESTION: 130

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: ([SHOW ANSWER](#))

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: 131

Which virtual warehouse configuration requires the maximum number of clusters to be equal to the minimum number of clusters?

- A. Standard scaling policy
- B. Economy scaling policy
- C. Auto-scale mode
- D. Maximized mode

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 132

What command will grant the SELECT privilege on all new tables created in the mydb.myschema schema to the read_only role?

- A. GRANT USAGE ON FUTURE SCHEMAS IN DATABASE mydb TO ROLE read_only;
- B. GRANT USAGE ON ALL SCHEMAS IN DATABASE mydb TO ROLE read_only;
- C. GRANT SELECT ON FUTURE TABLES IN SCHEMA mydb.myschema TO ROLE read_only;
- D. GRANT SELECT ON ALL TABLES IN SCHEMA mydb.myschema TO ROLE read_only;

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 133

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

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

Answer: ([SHOW ANSWER](#))

To change the existing file format type from CSV to JSON, the recommended way is to use the ALTER FILE FORMAT command with the SET TYPE=JSON clause. This alters the file format specification to use JSON instead of CSV.

References: Based on my internal knowledge as of 2021.

NEW QUESTION: 134

If Role_A owns Role_B, what permissions does Role_A have based on this relationship?

- A. Role_A automatically inherits all privileges assigned to Role_B.
- B. Role_A can use the privileges of Role_B, if Role_A is higher than Role_B in a defined hierarchy.
- C. Role_A can grant or revoke access to the role.
- D. Role_A can execute any tasks assigned to Role_B, without an explicit grant.

Answer: ([SHOW ANSWER](#))

The correct answer is C. Role_A can grant or revoke access to the role .

In Snowflake, the OWNERSHIP privilege on an object gives full control over that object. A role can own another role. If Role_A owns Role_B, then Role_A has control over Role_B, including the ability to manage grants involving that role.

Why C is correct:

Ownership of a role allows the owning role to manage the role. This includes granting the role to users or other roles and revoking those grants.

Why the other options are incorrect:

A). Ownership of a role does not automatically mean the owner inherits all privileges assigned to that role.

Privilege inheritance happens through role grants in the role hierarchy, not merely through ownership.

B). Role_A can use Role_B privileges only if Role_B is granted to Role_A in the role hierarchy. Ownership alone is not the same as inheritance.

D). A task requires appropriate task privileges, such as OPERATE or ownership, and account-level task execution privileges where applicable. Ownership of a role does not automatically allow executing tasks assigned to that role.

Official Snowflake documentation reference:

Snowflake documentation explains that roles are securable objects and that the OWNERSHIP privilege grants full control over an object. Role inheritance is separate and occurs when one role is granted to another role.

Reference: Snowflake Documentation - Access control overview; Snowflake Documentation - Role hierarchy and privilege inheritance; Snowflake Documentation - GRANT ROLE; Snowflake Documentation

- OWNERSHIP privilege; SnowPro Core Study Guide - Security and Access Control.

NEW QUESTION: 135

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

A. Clustering

B. Materialized Views

C. Zero-copy Cloning

D. Search Optimization Service

Answer: (SHOW ANSWER)

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

NEW QUESTION: 136

The bulk data load history that is available upon completion of the COPY statement is stored where and for how long?

A. In the metadata of the target table for 14 days

B. In the metadata of the pipe for 14 days

C. In the metadata of the target table for 64 days

D. In the metadata of the pipe for 64 days

Answer: (SHOW ANSWER)

The bulk data load history available after a COPY statement is stored in the metadata of the pipe and is retained for 64 days¹.

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 137

What Snowflake features are recommended to restrict unauthorized users from accessing Personal Identifiable Information (PII)? (Select TWO).

- A. Transient tables
- B. Secure views
- C. Data encryption
- D. Dynamic Data Masking
- E. Multi-Factor Authentication (MFA)

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 138

When unloading data, which file format preserves the data values for floating-point number columns?

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

Answer: ([SHOW ANSWER](#))

When unloading data, the Parquet file format is known for its efficiency in preserving the data values for floating-point number columns. Parquet is a columnar storage file format that offers high compression ratios and efficient data encoding schemes. It is especially effective for floating-point data, as it maintains high precision and supports efficient querying and analysis.

References:

Snowflake Documentation: Using the Parquet File Format for Unloading Data

NEW QUESTION: 139

Which statement describes Snowflake tables?

- A. Snowflake tables are logical representation of underlying physical data.
- B. Snowflake tables are the physical instantiation of data loaded into Snowflake.
- C. Snowflake tables require that clustering keys be defined to perform optimally.
- D. Snowflake tables are owned by a user.

Answer: ([SHOW ANSWER](#))

In Snowflake, tables represent a logical structure through which users interact with the stored data. The actual physical data is stored in micro-partitions managed by Snowflake, and the logical table structure provides the means by which SQL operations are mapped to this data. This architecture allows Snowflake to optimize storage and querying across its distributed, cloud-based data storage system.

References: Snowflake Documentation on Tables

NEW QUESTION: 140

What information does a Snowflake directory table hold?

- A. Metadata of all views in a specific schema
- B. Metadata of all tables in a specific schema

- C. Metadata of clustering keys of all tables in a specific schema
- D. File-level metadata of all data files in an internal or external stage

Answer: ([SHOW ANSWER](#))

The correct answer is D. File-level metadata of all data files in an internal or external stage .

A Snowflake directory table is associated with a stage and stores file-level metadata for files in that stage. It lets users query staged file information using SQL.

Why D is correct:

Directory tables hold metadata about staged files, such as file name, relative path, size, last modified time, and file URL information.

Example concept:

```
SELECT *
```

```
FROM DIRECTORY(@my_stage);
```

This returns metadata for files in the stage associated with the directory table.

Why the other options are incorrect:

- A). Directory tables do not store metadata about views in a schema.
- B). Directory tables do not store metadata about regular Snowflake tables in a schema.
- C). Directory tables do not store clustering key metadata.

Official Snowflake documentation reference:

Snowflake documentation describes directory tables as objects that store file-level metadata about staged files in cloud storage.

Reference: Snowflake Documentation - Directory tables; Snowflake Documentation - Stages; SnowPro Core Study Guide - Data Loading and Unloading.

NEW QUESTION: 141

Which property or parameter can be used to temporarily disable Multi-Factor Authentication (MFA) for a Snowflake user?

- A. ALLOW_CLIENT_MFA_CACHING
- B. EXT_AUTHN_DUO
- C. DISABLE_MFA
- D. MINS_TO_BYPASS_MFA

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 142

An external stage many_stage contains many directories including one, app_files that contains CSV files How can all the CSV files from this directory be moved into table my_table without scanning files that are not needed?

A. sql

```
COPY INTO my_table
```

```
FROM @many_stage
```

```
WHERE metadata$filename LIKE 'app_files/%'
```

```
FILE_FORMAT = (TYPE = CSV);
```

B. sql

```
COPY INTO my_table
FROM @many_stage
FILE_FORMAT = (TYPE = CSV)
PATTERN = 'app_files/*.csv';
```

C. sql

```
COPY INTO my_table
FROM @many_stage/app_files
FILE_FORMAT = (TYPE = CSV);
```

D. sql

```
COPY INTO my_table
FROM @many_stage
FILES = ('app_files/*.csv')
FILE_FORMAT = (TYPE = CSV);
```

E. sql

```
ALTER STAGE many_stage SET PATH = 'app_files';
COPY INTO my_table FROM @many_stage
FILE_FORMAT = (TYPE = CSV);
```

F. sql

```
COPY INTO my_table
FROM @many_stage
FILE_FORMAT = (TYPE = CSV)
FOLDER = 'app_files';
```

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 143

Which Snowflake layer is associated with virtual warehouses?

- A.** Cloud services
- B.** Query processing
- C.** Elastic memory
- D.** Database storage

Answer: ([SHOW ANSWER](#))

The layer of Snowflake 's architecture associated with virtual warehouses is the Query Processing layer.

Virtual warehouses in Snowflake are dedicated compute clusters that execute SQL queries against the stored data.

This layer is responsible for the entire query execution process, including parsing, optimization, and the actual computation. It operates independently of the storage layer, enabling Snowflake to scale compute and storage resources separately for efficiency and cost-effectiveness.

References:

Snowflake Documentation: Snowflake Architecture

NEW QUESTION: 144

A user has unloaded data from Snowflake to a stage

Which SQL command should be used to validate which data was loaded into the stage?

- A. list @file__stage
- B. show @file__stage
- C. view @file__stage
- D. verify @file__stage

Answer: ([SHOW ANSWER](#))

The `list` command in Snowflake is used to validate and display the list of files in a specified stage. When a user has unloaded data to a stage, running the `list @file__stage` command will show all the files that have been uploaded to that stage, allowing the user to verify the data that was unloaded.

References:

Snowflake Documentation on Stages

SnowPro Core Certification Study Guide

NEW QUESTION: 145

Which database objects can be shared using Secure Data Sharing? (Select TWO).

- A. Users
- B. Tables
- C. Roles
- D. Stages
- E. Dynamic tables

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 146

What are characteristics of the ownership privilege when it is granted on a regular Snowflake schema? (Select TWO).

- A. It is automatically granted to the role that creates a database object within the schema.
- B. It allows a role to manage grants on the schema.
- C. It can be transferred from one role to another for a specific schema.
- D. It grants the ability to query data from the schema.
- E. It must be granted to a role in order to alter warehouse settings.

Answer: ([SHOW ANSWER](#))

In Snowflake, the ownership privilege for a schema includes:

Automatic granting to the creator's role: The role that creates a database object within the schema automatically receives ownership of that object.

Ability to manage grants: Ownership enables the role to manage permissions and grants on the schema and its objects, allowing them to control access at the schema level.

Ownership does not directly confer query privileges or the ability to alter warehouse settings, nor is it transferable without specific privilege management actions.

NEW QUESTION: 147

What is the default size for a newly-provisioned Snowpark-optimized virtual warehouse ?

- A. X-Large

- B. Large
- C. Medium
- D. Small

Answer: ([SHOW ANSWER](#))

By default, Snowpark-optimized warehouses are provisioned with Medium size , unless specified otherwise. This provides balanced compute for handling Snowpark operations.

Reference:

Snowflake Docs: Snowpark-Optimized Warehouses

NEW QUESTION: 148

Which file function generates a Snowflake-hosted URL that must be authenticated when used?

- A. GET_STATE_LOCATION
- B. GET_PRESENT_URL
- C. BUILD_SCOPED_FILE_URL
- D. BUILD_STAGE_FILE_URL

Answer: ([SHOW ANSWER](#))

Purpose: The BUILD_STAGE_FILE_URL function generates a temporary, pre-signed URL that allows you to access a file within a Snowflake stage (internal or external). This URL requires authentication to use.

Key Points:

Security: The URL has a limited lifespan, enhancing security.

Use Cases: Sharing staged data with external tools or applications, or downloading it directly.

Snowflake Documentation (BUILD_STAGE_FILE_URL): https://docs.snowflake.com/en/sql-reference/functions/build_stage_file_url.html

NEW QUESTION: 149

A user needs to load Snowflake files from a stage that contains data from the past 5 years while using the least amount of operational overhead. The file names follow a specific format, such as files ending in _01.gz.

Which command should be used to load the data from 2024?

- A. Use the COPY INTO command to individually load the needed files.
- B. Use the COPY INTO command using the PATTERN parameter to define the needed files.
- C. Use the REMOVE command to eliminate unwanted files, then use the COPY INTO command.
- D. Use the GET command to download the needed files, then create a new stage and load the files onto the stage using the COPY INTO command.

Answer: ([SHOW ANSWER](#))

The correct answer is B. Use the COPY INTO command using the PATTERN parameter to define the needed files . The PATTERN parameter in the COPY INTO < table > command allows Snowflake to load only files whose names match a specified regular expression. This is useful when a stage contains many files and the required files can be identified by a naming pattern, such as year or suffix.

Why B is correct:

The user needs to load a subset of staged files from a 5-year collection and wants the least operational overhead.

Using PATTERN avoids manually listing files, removing files, or downloading and restaging data.

Example concept:

```
COPY INTO target_table
```

```
FROM @my_stage
```

```
PATTERN = '.*2024.*_01[.]gz ' ;
```

Why the other options are incorrect:

A). Individually loading files creates more manual work and higher operational overhead.

C). Removing unwanted files is risky and unnecessary. It also changes the staged file set.

D). Downloading files with GET, creating a new stage, and reloading files adds unnecessary operational complexity.

Official Snowflake documentation reference:

Snowflake documentation explains that the PATTERN option in COPY INTO < table > specifies a regular expression pattern used to match staged file names for loading.

Reference: Snowflake Documentation - COPY INTO < table > ; Snowflake Documentation - PATTERN copy option; SnowPro Core Study Guide - Data Loading and Unloading.

NEW QUESTION: 150

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 (LEAVE A REPLY)

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: 151

Which object can be shared from a database in an account with other Snowflake accounts using Secure Data Sharing?

A. Event table

B. Iceberg table

C. Temporary table

D. Stored procedure

Answer: (SHOW ANSWER)

The correct answer is B. Iceberg table .

Snowflake Secure Data Sharing allows providers to share supported database objects with other Snowflake accounts. Supported objects include tables, secure views, secure materialized views, secure UDFs, and certain other supported table types. Snowflake-managed Iceberg tables can be shared in supported configurations.

Why B is correct:

An Iceberg table is a supported table object in Snowflake and can be shared using Secure Data Sharing, subject to Snowflake's supported Iceberg table sharing requirements.

Why the other options are incorrect:

- A). Event tables are used to store logging and tracing event data. They are not the expected shareable object in this question.
- C). Temporary tables exist only within the session that created them and cannot be shared with other accounts.
- D). Stored procedures are not standard objects shared through Secure Data Sharing in the same way as tables and secure views.

Official Snowflake documentation reference:

Snowflake documentation lists supported objects for Secure Data Sharing and describes sharing supported table objects, including Iceberg tables where supported.

Reference: Snowflake Documentation - Secure Data Sharing; Shareable database objects; Iceberg tables; SnowPro Core Study Guide - Data Protection and Governance.

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (**1187** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 152

Which type of URL gives permanent access to files in cloud storage?

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 153

A sales table FCT_SALES has 100 million records.

The following Query was executed

```
SELECT COUNT (1) FROM FCT__SALES;
```

How did Snowflake fulfill this query?

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

Answer: ([SHOW ANSWER](#))

Snowflake is designed to optimize query performance by utilizing metadata for certain types of queries.

When executing a COUNT query, Snowflake can often fulfill the request by accessing metadata about the table's row count, rather than scanning the entire table or micro-partitions. This is particularly efficient for large tables like FCT_SALES with a significant number of records. The metadata layer maintains statistics about the table, including the

row count, which enables Snowflake to quickly return the result of a COUNT query without the need to perform a full scan.

References:

Snowflake Documentation on Metadata Management

SnowPro Core Certification StudyGuide

NEW QUESTION: 154

In Snowpipe, what is the default behavior of the ON_ERROR copy option?

- A. CONTINUE
- B. SKIP_FILE
- C. SKIP_FILE_ < num >
- D. ABORT_STATEMENT

Answer: ([SHOW ANSWER](#))

The correct answer is B. SKIP_FILE .

Important correction:

The provided answer A. CONTINUE is not correct for Snowpipe.

For regular bulk loading with COPY INTO < table > , the default ON_ERROR behavior is commonly associated with ABORT_STATEMENT. However, Snowpipe has different behavior. In Snowpipe, the default ON_ERROR option is SKIP_FILE .

Why B is correct:

With ON_ERROR = SKIP_FILE, if Snowpipe encounters an error while loading a file, the file is skipped instead of causing the entire continuous loading process to fail.

Why the other options are incorrect:

- A). CONTINUE continues loading and skips individual error rows, but it is not the default Snowpipe behavior.
- C). SKIP_FILE_ < num > skips a file after a specified number of errors, but it is not the default.
- D). ABORT_STATEMENT is associated with standard COPY INTO < table > behavior, not the default Snowpipe behavior.

Official Snowflake documentation reference:

Snowflake documentation for Snowpipe explains that the default ON_ERROR copy option for Snowpipe is SKIP_FILE.

Reference: Snowflake Documentation - Snowpipe; Snowflake Documentation - CREATE PIPE; Snowflake

Documentation - COPY INTO < table > copy options; SnowPro Core Study Guide - Data Loading and Unloading.

NEW QUESTION: 155

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 ([LEAVE A REPLY](#))

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: 156

Dynamic Data Masking can be used with which Snowflake objects? (Select TWO)

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

Answer: ([SHOW ANSWER](#))

Dynamic Data Masking (DDM) can be applied to:

- * Tables
- * Views

It cannot be applied to materialized views or external tables as those do not support the runtime evaluation of masking policies.

Reference:

Snowflake Docs: Dynamic Data Masking

NEW QUESTION: 157

What costs are associated with sharing Snowflake data with a reader account ?

- A. The data provider incurs costs for the data storage costs, and the data consumer incurs compute costs.
- B. The data provider incurs costs for compute, and the data consumer incurs costs for data storage.
- C. The data provider does not incur costs when sharing data; the data consumer incurs costs for data storage and compute.
- D. The data consumer does not incur costs when accessing shared data; the data provider incurs costs for data storage and compute.

Answer: ([SHOW ANSWER](#))

With reader accounts , the provider pays for:

- * Storage of shared data
- * Compute used by the reader account

The consumer does not pay , as the infrastructure is provisioned and managed by the provider.

Reference:

Snowflake Docs: Reader Accounts and Costs

NEW QUESTION: 158

Which type of charts are supported by Snowsight? {Select TWO}

- A. Line charts
- B. Pie charts
- C. Gantt charts

D. Scatterplots

E. Flowcharts

Answer: B,D ([LEAVE A REPLY](#))

NEW QUESTION: 159

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

A. @

B. &

C. \$

D. #

Answer: C ([LEAVE A REPLY](#))

VeryComprehensiveExplanation=InSnowflake,variablesaredenotedbyadollarsign(). Variables 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: 160

What is the outcome of querying a VARIANT column using an invalid path?

A. An empty string

B. NULL

C. An invalid identifier error

D. Undefined data

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 161

Which Snowflake feature or service is primarily used for managing and monitoring data and user activities?

A. Snowflake Marketplace

B. SnowSQL

C. Streamlit

D. Snowsight

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 162

Which command can be used to list all network policies available in an account?

A. DESCRIBE SESSION POLICY

B. DESCRIBE NETWORK POLICY

C. SHOW SESSION POLICIES

D. SHOW NETWORK POLICIES

Answer: ([SHOW ANSWER](#)**)**

To list all network policies available in an account, the correct command isSHOW NETWORK POLICIES.

Network policies in Snowflake are used to define and enforce rules for how users can connect to Snowflake, including IP whitelisting and other connection requirements. The `SHOW NETWORK POLICIES` command provides a list of all network policies defined within the account, along with their details.

The `DESCRIBE SESSION POLICY` and `DESCRIBE NETWORK POLICY` commands do not exist in Snowflake SQL syntax. The `SHOW SESSION POLICIES` command is also incorrect, as it does not pertain to the correct naming convention used by Snowflake for network policy management.

Using `SHOW NETWORK POLICIES` without any additional parameters will display all network policies in the account, which is useful for administrators to review and manage the security configurations pertaining to network access.

Reference: Snowflake Documentation on Network Policies (<https://docs.snowflake.com/en/sql-reference/sql/show-network-policies.html>)

NEW QUESTION: 163

A user wants to upload a file to an internal Snowflake stage using a put command.

Which tools and or connectors could be used to execute this command? (Select TWO).

- A. SnowCD
- B. SnowSQL
- C. SQL API
- D. Python connector
- E. Snowsight worksheets

Answer: ([SHOW ANSWER](#))

To upload a file to an internal Snowflake stage using a PUT command, you can use:

B). SnowSQL: SnowSQL, the command-line client for Snowflake, supports the PUT command, allowing users to upload files directly to Snowflake stages from their local file systems.

E). Snowsight worksheets: Snowsight, the web interface for Snowflake, provides a user-friendly environment for executing SQL commands, including the PUT command, through its interactive worksheets.

References:

Snowflake Documentation: Loading Data into Snowflake using SnowSQL

Snowflake Documentation: Using Snowsight

NEW QUESTION: 164

When tracking data storage metrics, how can a user identify deleted bytes that are still accruing storage charges?

- A. Query the `TABLE_STORAGE_METRICS` view in the `INFORMATION_SCHEMA`.
- B. Query the `DATABASE_STORAGE_USAGE_HISTORY` view in `ACCOUNT_USAGE`.
- C. Use the `DATABASE_STORAGE_USAGE_HISTORY` table function.
- D. Query the `STORAGE_USAGE` view in `ACCOUNT_USAGE`.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 165

Which solution improves the performance of point lookup queries that return a small number of rows from large tables using highly selective filters?

- A. Automatic clustering

- B. Materialized views
- C. Query acceleration service
- D. Search optimization service

Answer: ([SHOW ANSWER](#))

The search optimization service improves the performance of point lookup queries on large tables by using selective filters to quickly return a small number of rows. It creates an optimized data structure that helps in pruning the micro-partitions that do not contain the queried values³. References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 166

Why would a Snowflake user decide to use a materialized view instead of a regular view?

- A. The base tables do not change frequently.
- B. The results of the view change often.
- C. The query is not resource intensive.
- D. The query results are not used frequently.

Answer: ([SHOW ANSWER](#))

A Snowflake user would decide to use a materialized view instead of a regular view primarily when the base tables do not change frequently. Materialized views store the result of the view query and update it as the underlying data changes, making them ideal for situations where the data is relatively static and query performance is critical. By precomputing and storing the query results, materialized views can significantly reduce query execution times for complex aggregations, joins, and calculations.

References:

Snowflake Documentation: Materialized Views

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 167

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: ([SHOW ANSWER](#))

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: 168

A user notices that query performance has decreased. The virtual warehouses appear to be properly sized and are functioning correctly. The user suspects the issue lies within the services responsible for query planning and access validation. Which components of Snowflake architecture would be involved in diagnosing or optimizing this issue? Select TWO.

- A. The services handling authentication and access control.
- B. The services managing data loading and ETL operations.
- C. The services controlling warehouse scaling and suspension.
- D. The services managing data micro-partitioning and file storage.
- E. The services responsible for query parsing and optimization.

Answer: ([SHOW ANSWER](#))

Snowflake's Cloud Services layer handles authentication, access control, metadata management, query parsing, and query optimization. Since the issue is suspected to involve query planning and access validation, the relevant components are access control and query parsing/optimization.

Reference: Snowflake Documentation - Snowflake architecture; Cloud Services layer.

NEW QUESTION: 169

A query containing a WHERE clause is running longer than expected. The Query Profile shows that all micro-partitions being scanned. How should this query be optimized?

- A. Create a view on the table.
- B. Add a clustering key to the table.
- C. Add a limit clause to the query.
- D. Add a Dynamic Data Masking policy to the table.

Answer: ([SHOW ANSWER](#))

When a query containing a WHERE clause is running longer than expected, and the Query Profile shows that all micro-partitions are being scanned, the query can be optimized by adding a clustering key to the table.

Understanding Micro-Partitioning in Snowflake:

Snowflake automatically partitions tables into micro-partitions for efficient storage and query performance.

Each micro-partition contains metadata about the range of values it holds, which helps in pruning irrelevant partitions during query execution.

Role of Clustering Keys:

A clustering key defines how data in a table is organized within micro-partitions.

By specifying a clustering key, you can control the physical layout of data, ensuring that related rows are stored together.

This organization improves query performance by reducing the number of micro-partitions that need to be scanned.

Optimizing Queries with Clustering Keys:

Adding a clustering key based on columns frequently used in WHERE clauses helps Snowflake quickly locate and scan relevant micro-partitions.

This minimizes the amount of data scanned and reduces query execution time.

Example:

```
ALTER TABLE my_table CLUSTER BY (column1, column2);
```

This command adds a clustering key to my_table using column1 and column2.

Future queries that filter on these columns will benefit from improved performance.

Benefits:

Reduced query execution time: Fewer micro-partitions need to be scanned.

Improved resource utilization: More efficient data retrieval leads to lower compute costs.

Snowflake Documentation: Clustering Keys

Snowflake Documentation: Query Profile

NEW QUESTION: 170

What is a feature of column-level security in Snowflake?

- A. Role access policies
- B. Network policies
- C. Internal tokenization
- D. External tokenization

Answer: ([SHOW ANSWER](#))

Column-level security in Snowflake is implemented through Role Access Policies. These policies allow administrators to control access to specific columns of a table or view based on the role of the user accessing the data. By applying a role access policy to a column, administrators can ensure that sensitive information remains secure, and only users with the appropriate roles can view or query the data in that column. This feature enhances the security model by providing fine-grained access control at the column level.

References:

Snowflake Documentation: Implementing Column-level Security

NEW QUESTION: 171

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 172

What is the least-privileged database role needed to view the definition of a secure materialized view?

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

Answer: ([SHOW ANSWER](#))

The USAGE_VIEWER role is the minimum required to view metadata and definitions of governance- enabled or secure objects like secure materialized views .

Reference:

Snowflake Docs: Governance Roles

NEW QUESTION: 173

Which system_defined, read-only view display information on column lineage that specifies how data flows from source to target in a SQL write operation?

- A. ACCESS_HISTORY
- B. LOAD_HOSTORY
- C. QUERY_HISTORY
- D. COPY_HISTORY

Answer: A ([LEAVE A REPLY](#))

In Snowflake, the system-defined, read-only view that displays information on column lineage, which specifies how data flows from source to target in a SQL write operation, isACCESS_HISTORY. This view is instrumental in auditing and analyzing data access patterns, as it provides detailed insights into how and from where the data is being accessed and manipulated within Snowflake.

Reference to Snowflake documentation on ACCESS_HISTORY:

Using Access History to Audit Data Access

NEW QUESTION: 174

Which of the following describes a Snowflake stored procedure?

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 175

Which statement describes how Snowflake supports reader accounts?

- A. A reader account can consume data from the provider account that created it and combine it with its own data.
- B. A consumer needs to become a licensed Snowflake customer as data sharing is only supported between Snowflake accounts.

C. The users in a reader account can query data that has been shared with the reader account and can perform DML tasks.

D. The SHOW MANAGED ACCOUNTS command will view all the reader accounts that have been created for an account.

Answer: ([SHOW ANSWER](#))

Snowflake supports reader accounts, which are a type of account that allows data providers to share data with consumers who are not Snowflake customers. However, for data sharing to occur, the consumer needs to become a licensed Snowflake customer because data sharing is only supported between Snowflake accounts.

References: Introduction to Secure Data Sharing | Snowflake Documentation2.

NEW QUESTION: 176

For which use case will transient tables be most cost-effective?

A. The table data does not require Fail-safe.

B. The table data requires long-term storage.

C. The table data is frequently accessed and modified.

D. The table data is frequently shared with external accounts.

Answer: ([SHOW ANSWER](#))

The correct answer is A. The table data does not require Fail-safe .

Transient tables are designed for data that does not need the same level of data recovery protection as permanent tables. Unlike permanent tables, transient tables do not have a Fail-safe period.

Why A is correct:

Fail-safe storage can add cost because Snowflake retains historical data for recovery by Snowflake support after the Time Travel retention period ends. Since transient tables do not have Fail-safe, they can reduce storage costs for data that does not require this additional protection.

Common use cases include:

Temporary ETL data

Staging data

Intermediate transformation data

Reproducible data

Non-critical working data

Why the other options are incorrect:

B). Long-term critical storage is usually better suited for permanent tables because they support Fail-safe.

C). Frequent access and modification alone does not make a transient table the best cost choice. The key factor is whether Fail-safe is needed.

D). Sharing data with external accounts is not the primary cost-effective use case for transient tables.

Official Snowflake documentation reference:

Snowflake documentation explains that transient tables persist until explicitly dropped, support Time Travel with limited retention, and do not have a Fail-safe period.

Reference: Snowflake Documentation - Transient tables; Snowflake Documentation - Understanding and viewing Fail-safe; SnowPro Core Study Guide - Snowflake Architecture.

NEW QUESTION: 177

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: ([SHOW ANSWER](#))

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: 178

What is the default format-type COMPRESSION option when unloading CSV data from Snowflake using the COPY INTO < location > command?

- A. GZIP
- B. AUTO
- C. NONE
- D. DEFLATE

Answer: ([SHOW ANSWER](#))

The correct answer is A. GZIP .

When unloading data from Snowflake to staged files using COPY INTO < location > , the default compression for CSV output is GZIP unless a different compression option is specified.

Why A is correct:

For CSV unloads, Snowflake compresses output files using GZIP by default.

Example:

```
COPY INTO @my_stage/output/  
FROM my_table  
FILE_FORMAT = (TYPE = CSV);
```

If no compression is specified, the unloaded CSV files are compressed with GZIP by default.

Why the other options are incorrect:

- B). AUTO is commonly used when loading data so Snowflake can automatically detect compression. It is not the default unload compression for CSV.
- C). NONE disables compression, but it is not the default.
- D). DEFLATE is a supported compression type, but it is not the default for CSV unload.

Official Snowflake documentation reference:

Snowflake documentation for CSV file format options explains that the default COMPRESSION value for unloading CSV data is GZIP.

Reference: Snowflake Documentation - COPY INTO < location > ; Snowflake Documentation - CSV file format options; SnowPro Core Study Guide - Data Loading and Unloading.

NEW QUESTION: 179

Which statement accurately describes how a virtual warehouse functions?

- A. Increasing the size of a virtual warehouse will always improve data loading performance.
- B. Each virtual warehouse is an independent compute cluster that shares compute resources with other warehouses.
- C. Each virtual warehouse is a compute cluster composed of multiple compute nodes allocated by Snowflake from a cloud provider.
- D. All virtual warehouses share the same compute resources so performance degradation of one warehouse can significantly affect all the other warehouses.

Answer: ([SHOW ANSWER](#))

A virtual warehouse in Snowflake is an independent compute cluster that performs data processing tasks such as executing SQL queries. Each virtual warehouse is dynamically allocated by Snowflake from the cloud provider 's resources and does not share compute resources with other warehouses. This architecture ensures that the performance of one warehouse does not impact the performance of another. Adjusting the size of a virtual warehouse affects its computational power by increasing or decreasing the number of compute nodes, which can improve the performance of data processing tasks depending on the workload.

References:

Snowflake Documentation: Understanding Virtual Warehouses

NEW QUESTION: 180

Which Snowflake privilege is required on a pipe object to pause or resume pipes?

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

Answer: ([SHOW ANSWER](#))

OPERATE. In Snowflake, to pause or resume a pipe, the OPERATE privilege is required on the pipe object.

The OPERATE privilege allows users to perform operational tasks on specific objects such as pipes, tasks, and streams. Specifically, for a pipe, the OPERATE privilege enables the user to execute the ALTER PIPE ... SET PIPE_EXECUTION_PAUSED=TRUE or ALTER PIPE ... SET PIPE_EXECUTION_PAUSED=FALSE commands, which are used to pause or resume the pipe, respectively.

Here 's a step-by-step explanation and reference:

Understanding Pipe Operations: Pipes in Snowflake are used for continuous data loading from staging areas into Snowflake tables. Managing pipes involves operations such as creating, monitoring, pausing, and resuming.

Privileges for Pipe Operations: The OPERATE privilege is essential for pausing and resuming pipes. This privilege is more specific than general object access privileges like SELECT or USAGE and is tailored for operational control.

Granting the OPERATE Privilege: To grant the OPERATE privilege on a pipe, an administrator or a user with the necessary grants can execute the SQL command:

```
GRANT OPERATE ON PIPE < pipe_name > TO ROLE < role_name > ;
```

Pausing and Resuming Pipes: Once the OPERATE privilege is granted, the user or role can pause the pipe using:

```
ALTERPIPE < pipe_name > SETPIPE_EXECUTION_PAUSED=TRUE;
```

To resume the pipe, they use:

```
ALTERPIPE < pipe_name > SETPIPE_EXECUTION_PAUSED=FALSE;
```

Reference: For more information on pipe operations and privileges, you can refer to Snowflake 's official documentation on pipes and security access control: <https://docs.snowflake.com/en/sql-reference/sql/grant-privilege.html> and <https://docs.snowflake.com/en/user-guide/data-pipelines-intro.html#managing-pipes>

NEW QUESTION: 181

What does Snowflake recommend for a user assigned the ACCOUNTADMIN role?

- A. The ACCOUNTADMIN role should be set as the user 's default role.
- B. The user should use federated authentication instead of a password
- C. The user should be required to use Multi-Factor Authentication (MFA).
- D. There should be just one user with the ACCOUNTADMIN role in each Snowflake account.

Answer: ([SHOW ANSWER](#))

For users assigned the ACCOUNTADMIN role, Snowflake recommends enforcing Multi-Factor Authentication (MFA) to enhance security. The ACCOUNTADMIN role has extensive permissions, making it crucial to secure accounts held by such users against unauthorized access. MFA adds an additional layer of security by requiring a second form of verification beyond just the username and password, significantly reducing the risk of account compromise.

References: Snowflake Security Best Practices

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here: <https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (**1187** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 182

What happens to event notifications that are received when an automated Snowpipe is paused?

- A. They enter a retention period of 14 days, after which they are dropped from the internal metadata.
- B. They are permanently stored in the internal metadata and processed once the pipe is resumed.
- C. They are stored indefinitely in the internal metadata until the pipe is manually resumed.
- D. They are immediately dropped from the internal metadata.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 183

By default, which tasks can a user with the ORGADMIN role perform? Select TWO.

- A. Create a share.
- B. Enable replication.

- C. Create an account.
- D. Create a resource monitor.
- E. Create a reader account.

Answer: ([SHOW ANSWER](#))

The correct answers are B. Enable replication and C. Create an account .

The ORGADMIN role is an organization-level role in Snowflake. It is used to manage organization-wide operations, including account management and enabling cross-account or cross-region capabilities such as replication.

Why B is correct:

ORGADMIN can enable replication for the organization. Replication is an organization-level capability because it can involve multiple accounts and regions.

Why C is correct:

ORGADMIN can create Snowflake accounts within the organization.

Why the other options are incorrect:

- A). Creating a share is an account-level data sharing task, not a default organization-level ORGADMIN task.
- D). Creating a resource monitor is an account-level resource management task, typically handled by account-level administrative roles such as ACCOUNTADMIN.
- E). Creating a reader account is an account-level secure data sharing task, not the primary default ORGADMIN responsibility in this answer set.

Official Snowflake documentation reference:

Snowflake documentation describes ORGADMIN as the organization administrator role used to manage organization-level operations, including creating accounts and enabling replication.

Reference: Snowflake Documentation - Organization administrator role; Snowflake Documentation - Creating accounts; Snowflake Documentation - Replication and failover; SnowPro Core Study Guide - Snowflake Account and Resource Management.

NEW QUESTION: 184

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 185

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

- A. External Tokenization

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 186

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

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

Answer: ([SHOW ANSWER](#))

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

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

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

NEW QUESTION: 187

Which Snowflake object helps evaluate virtual warehouse performance impacted by query queuing?

- A. Resource monitor
- B. Account_usage.query_history
- C. Information_schema.warehouse_load_history
- D. Information schema.warehouse metering history

Answer: ([SHOW ANSWER](#))

The Snowflake object that helps evaluate virtual warehouse performance impacted by query queuing is the Information_schema.warehouse_load_history. This view provides historical data about the load on a warehouse, including the average number of queries that were running or queued within a specific interval, which can be used to assess performance and identify potential issues with query queuing³.

NEW QUESTION: 188

What is SnowSQL used for?

- A. Developing dashboards
- B. Loading and unloading data
- C. Developing Snowflake Native Apps
- D. Monitoring data transformation activity

Answer: (SHOW ANSWER)

The correct answer is B. Loading and unloading data .

SnowSQL is Snowflake's command-line client. It is used to connect to Snowflake, execute SQL statements, run SQL scripts, and perform data loading and unloading operations.

Why B is correct:

SnowSQL supports data loading and unloading workflows. For example, users can use SnowSQL to run PUT, GET, and COPY commands.

Example loading workflow:

```
PUT file:///local/path/data.csv @my_stage;
```

```
COPY INTO my_table
```

```
FROM @my_stage
```

```
FILE_FORMAT = (TYPE = CSV);
```

Example unloading workflow:

```
COPY INTO @my_stage/output/
```

```
FROM my_table;
```

```
GET @my_stage/output/ file:///local/path/output/;
```

Why the other options are incorrect:

A). Dashboards are usually developed in BI tools or applications, not SnowSQL.

C). Snowflake Native Apps are developed and managed using Snowflake Native App Framework tools and Snowflake CLI, not primarily SnowSQL.

D). Monitoring transformation activity is usually done through Snowflake views, query history, task history, event tables, or monitoring features, not as the primary purpose of SnowSQL.

Official Snowflake documentation reference:

Snowflake documentation describes SnowSQL as the command-line client for connecting to Snowflake and executing SQL queries, including data loading and unloading commands.

Reference: Snowflake Documentation - SnowSQL; Snowflake Documentation - Loading data using SnowSQL;

Snowflake Documentation - PUT, GET, and COPY INTO; SnowPro Core Study Guide - Data Loading and Unloading.

NEW QUESTION: 189

Which user preferences can be set for a user profile in Snowsight? (Select TWO).

A. Multi-Factor Authentication (MFA)

B. Default database

C. Default schema

D. Notification

E. Username

Answer: (SHOW ANSWER)

In Snowsight, Snowflake 's web interface, user preferences can be customized to enhance the user experience.

Among these preferences, users can set a default database and default schema. These settings streamline the user experience by automatically selecting the specified database and schema when the user initiates a new session or query, reducing the need to manually specify these parameters for each operation.

This feature is particularly useful for users who frequently work within a specific database or schema context.

References: Snowflake Documentation on Snowsight User Preferences

NEW QUESTION: 190

Which ACCOUNT_USAGE view can be used to identify the masking policy assigned to an object?

- A. TAG_REFERENCES
- B. OBJECT_DEPENDENCIES
- C. POLICY_REFERENCES
- D. ACCESS_HISTORY

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 191

How does Snowflake ensure the integrity and consistency of data stored in tables?

- A. Row-level encryption
- B. Dynamic Data Masking
- C. Time Travel
- D. Constraints

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 192

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 193

Which privilege grants the ability to set a column-level security masking policy on a table or view column?

- A. APPLY
- B. CREATE
- C. SET
- D. MODIFY

Answer: ([SHOW ANSWER](#))

In Snowflake, the APPLY privilege is required to set a masking policy on a table or view column. This privilege allows a user to associate a masking policy with a specific column, thereby controlling how data in that column is masked or hidden.

Create a Masking Policy: Define a masking policy using the CREATE MASKING POLICY command.

Grant APPLY Privilege: Grant the APPLY privilege on the masking policy to the relevant roles or users.

Apply the Masking Policy: Use the ALTER TABLE or ALTER VIEW command to apply the masking policy to a column.

References:

Snowflake Documentation: Data Masking

Snowflake Documentation: CREATE MASKING POLICY

Snowflake Documentation: Privileges for Masking Policies

NEW QUESTION: 194

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

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

Answer: B,C ([LEAVE A REPLY](#))

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

NEW QUESTION: 195

Which type of role can be granted to a share?

- A. Account role
- B. Custom role
- C. Database role
- D. Secondary role

Answer: ([SHOW ANSWER](#)**)**

In Snowflake, shares are used to share data between Snowflake accounts. When creating a share, it is possible to grant access to the share to roles within the Snowflake account that is creating the share. The type of role that can be granted to a share is a Custom role. Custom roles are user-defined roles that account administrators can create to manage access control in a more granular way. Unlike predefined roles such as ACCOUNTADMIN or SYSADMIN, custom roles can be tailored with specific privileges to meet the security and access requirements of different groups within an organization.

Granting a custom role access to a share enables users associated with that role to access the shared data if the share is received by another Snowflake account. It is important to carefully manage the privileges granted to custom roles to ensure that data sharing aligns with organizational policies and data governance standards.

References:

Snowflake Documentation on Shares: Shares

Snowflake Documentation on Roles: Access Control

NEW QUESTION: 196

What information does the Query Profile provide?

- A. Graphical representation of the data model
- B. Statistics for each component of the processing plan

- C. Detailed Information about the database schema
- D. Real-time monitoring of the database operations

Answer: ([SHOW ANSWER](#))

The Query Profile in Snowflake provides a graphical representation and statistics for each component of the query's execution plan. This includes details such as the execution time, the number of rows processed, and the amount of data scanned for each operation within the query. The Query Profile is a crucial tool for understanding and optimizing the performance of queries, as it helps identify potential bottlenecks and inefficiencies.

References:

Snowflake Documentation: Understanding the Query Profile

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (**1187** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 197

How should a Snowflake user configure a virtual warehouse to be in "Maximized mode"?

- A. Set the WAREHOUSES_SIZE to 6XL.
- B. Set the STATEMENT_TIMEOUT_1M_SECONDS to 0.
- C. Set the MAX_CONCURRENCY_LEVEL to a value of 12 or larger.
- D. Set the same value for both MIN_CLUSTER_COUNT and MAX_CLUSTER_COUNT.

Answer: ([SHOW ANSWER](#))

In Snowflake, configuring a virtual warehouse to be in a "Maximized" mode implies maximizing the resources allocated to the warehouse for its duration. This is done to ensure that the warehouse has a consistent amount of compute resources available, enhancing performance for workloads that require a high level of parallel processing or for handling high query volumes.

To configure a virtual warehouse in maximized mode, you should set the same value for both MIN_CLUSTER_COUNT and MAX_CLUSTER_COUNT. This configuration ensures that the warehouse operates with a fixed number of clusters, thereby providing a stable and maximized level of compute resources.

Reference to Snowflake documentation on warehouse sizing and scaling:

Warehouse Sizing and Scaling

Understanding Warehouses

NEW QUESTION: 198

In Snowflake, the use of federated authentication enables which Single Sign-On (SSO) workflow activities? (Select TWO).

- A. Authorizing users
- B. Initiating user sessions

- C. Logging into Snowflake
- D. Logging out of Snowflake
- E. Performing role authentication

Answer: ([SHOW ANSWER](#))

Federated authentication in Snowflake allows users to use their organizational credentials to log in to Snowflake, leveraging Single Sign-On (SSO). The key activities enabled by this setup include:

- B). Initiating user sessions: Federated authentication streamlines the process of starting a user session in Snowflake by using the existing authentication mechanisms of an organization.
- C). Logging into Snowflake: It simplifies the login process, allowing users to authenticate with their organization ' s identity provider instead of managing separate credentials for Snowflake.

References:

Snowflake Documentation: Configuring Federated Authentication

NEW QUESTION: 199

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 200

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

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

Answer: ([SHOW ANSWER](#))

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

References:

[COF-C03] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Machine Learning & Data Science Partners

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

NEW QUESTION: 201

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

- A. Enable the query acceleration service.

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

Answer: C ([LEAVE A REPLY](#))

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

NEW QUESTION: 202

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

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

Answer: ([SHOW ANSWER](#)**)**

The Snowflake feature that allows for small volumes of data to be continuously loaded into Snowflake and incrementally made available for analysis is Snowpipe. Snowpipe is designed for near-real-time data loading, enabling data to be loaded as soon as it's available in the storage layer³

NEW QUESTION: 203

What is included in the calculation of the execution time reported in a Query Profile? (Select TWO).

- A. Spilling
- B. Processing
- C. Pruning
- D. Caching
- E. Local disk IO

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 204

A Snowflake user accidentally deleted a table. The table no longer exists, but the session is within the data retention period. How can the table be restored using the LEAST amount of operational overhead?

- A. Clone the table schema as it existed before the table was dropped.
- B. Clone the database as it existed before the table was dropped.
- C. Recreate the table and reload the data.
- D. Run the UNDROP command against the table.

Answer: ([SHOW ANSWER](#)**)**

In Snowflake, if a table is accidentally dropped but still within the data retention period (also known as " Time Travel "), the simplest and most efficient recovery method is the UNDROP command. This command restores the deleted table to its previous state with minimal operational effort. Since Snowflake retains dropped table data for a specific retention

period (up to 90 days for the Enterprise edition), UNDROP can quickly recover the table without the need for complex cloning or data reloading processes, making it ideal for accidental deletions.

NEW QUESTION: 205

What Snowflake recommendation is designed to ensure that staged data is only loaded once " ?

- A. Partitioning staged data files
- B. Loading only the most recently-staged data files
- C. Removing data files after loading
- D. Identifying and removing duplicates after each data load

Answer: ([SHOW ANSWER](#))

Snowflake recommends removing data files from the staging area after they have been loaded into the target table. This practice ensures that the data is only loaded once and prevents accidental reloading of the same data. By removing the files, you eliminate the risk of duplicate data loads.

Stage the Data: Upload the data files to a Snowflake stage (internal or external).

Load the Data: Use the COPY INTO command to load the data from the stage into the Snowflake table.

Remove the Data Files: After successfully loading the data, remove the data files from the stage using the REMOVE command.

References:

Snowflake Documentation: Loading Data into Snowflake

Snowflake Documentation: Staging Data Files

Snowflake Documentation: COPY INTO Command

NEW QUESTION: 206

Snowflake ' s access control framework combines which models for securing data? (Select TWO).

- A. Attribute-based Access Control (ABAC)
- B. Discretionary Access Control (DAC)
- C. Access Control List (ACL)
- D. Role-based Access Control (RBAC)
- E. Rule-based Access Control (RuBAC)

Answer: ([SHOW ANSWER](#))

Snowflake ' s access control framework utilizes a combination of Discretionary Access Control (DAC) and Role-based Access Control (RBAC). DAC in Snowflake allows the object owner to grant access privileges to other roles. RBAC involves assigning roles to users and then granting privileges to those roles. Through roles, Snowflake manages which users have access to specific objects and what actions they can perform, which is central to security and governance in the Snowflake environment.

References: Snowflake Documentation on Access Control,

NEW QUESTION: 207

If a network policy is applied at both the account and user levels, which policy takes precedence?

- A. The user-level policy
- B. The account-level policy

- C. The policy that was applied first
- D. The most-restrictive policy

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 208

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

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

Answer: ([SHOW ANSWER](#))

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

References:

[COF-C03] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Masking Policies

NEW QUESTION: 209

What is the impact of removing some of the privileges granted to a role?

- A. The role will retain all privileges until a new session is started.
- B. The role will lose access to objects that are not granted to the updated role.
- C. The user will have to reinstate the role using the USE ROLE command after the change.
- D. Any users who are actively using the role will have their current session terminated.

Answer: ([SHOW ANSWER](#))

The correct answer is B. The role will lose access to objects that are not granted to the updated role .

Snowflake access control is based on privileges granted to roles. Users receive access to objects through the active role and any inherited roles in the role hierarchy. If privileges are revoked from a role, that role no longer has those privileges unless they are still inherited through another granted role.

Why B is correct:

When privileges are removed from a role, users using that role lose the ability to perform actions that depended on the revoked privileges. Access is controlled by the current set of grants assigned to the role.

Why the other options are incorrect:

- A). A role does not retain revoked privileges until a new session starts. Revoked privileges affect access according to Snowflake access control rules.
- C). The user does not need to reissue USE ROLE just because privileges were revoked.
- D). Snowflake does not terminate active sessions simply because a privilege is removed from a role.

Official Snowflake documentation reference:

Snowflake documentation explains that privileges are granted to roles, roles are granted to users or other roles, and revoking privileges removes the ability to perform the associated actions unless access exists through another role path.

Reference: Snowflake Documentation - Access control overview; Snowflake Documentation - GRANT privileges; Snowflake Documentation - REVOKE privileges; SnowPro Core Study Guide - Security and Access Control.

NEW QUESTION: 210

A Snowflake user is actively logged into Snowflake when a user-level network policy is assigned to that user.

Which Snowflake do it the user 's IP address does not match the user-level network policy rules?

- A. Allow the user to continue until the session or login token expires.
- B. Deactivate the network policy.
- C. Log the user out.
- D. Prevent the user from executing additional queries.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 211

What situation is likely to cause data spillage when a query is run?

- A. When the number of queries exceeds the max_concurrency_level parameter setting.
- B. When the query contains multiple filters and no data is returned.
- C. When a virtual warehouse runs out of memory while executing the query.
- D. When running queries that exceed the statement_timecut_in_secnds parameter setting.

Answer: ([SHOW ANSWER](#))

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 212

Which function returns the URL of a stage using the stage name as the input?

- A. BUILD_STAGE_FILE_URL
- B. BUILD_SCOPED_FILE_URL
- C. GET_PREIGNED_URL
- D. GET STAGE LOCATION

Answer: ([SHOW ANSWER](#))

The function in Snowflake that returns the URL of a stage using the stage name as the input is C.

GET_PREIGNED_URL. This function generates a pre-signed URL for a specific file in a stage, enabling secure, temporary access to that file without requiring Snowflake credentials. While the function is primarily used for accessing files in external stages, such as Amazon S3 buckets, it is instrumental in scenarios requiring direct, secure file access for a limited time.

It ' s important to note that as of my last update, Snowflake ' s documentation does not specifically list a function named GET_PREIGNED_URL for directly obtaining a stage ' s URL by its name. The description aligns closely with functionality available in cloud storage services (e.g., AWS S3 ' s presigned URLs) which can be used in conjunction with Snowflake stages for secure, temporary access to files. For direct interaction with stages and their files, Snowflake offers various functions and commands, but the exact match for generating a presigned URL through a simple function call may vary or require leveraging external cloud services APIs in addition to Snowflake ' s capabilities.

References:

Snowflake Documentation and cloud services (AWS, Azure, GCP) documentation on presigned URLs and stage interactions.

NEW QUESTION: 213

What privilege should a user be granted to change permissions for new objects in a managed access schema?

- A. Grant the OWNERSHIP privilege on the schema.
- B. Grant the OWNERSHIP privilege on the database.
- C. Grant the MANAGE GRANTS global privilege.
- D. Grant ALL privileges on the schema.

Answer: (SHOW ANSWER)

To change permissions for new objects in a managed access schema, a user should be granted the MANAGE GRANTS global privilege. This privilege allows the user to manage access control through grants on all securable objects within Snowflake2. References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 214

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

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

Answer: (SHOW ANSWER)

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

NEW QUESTION: 215

What is the purpose of an External Function?

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

Answer: ([SHOW ANSWER](#))

The purpose of an External Function in Snowflake is to call code that executes outside of the Snowflake environment. This allows Snowflake to interact with external services and leverage functionalities that are not natively available within Snowflake, such as calling APIs or running custom code hosted on cloud services³.

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

NEW QUESTION: 216

How is Single Sign-On (SSO) authentication used in Snowflake?

- A.** SSO is an integrated Snowflake feature powered by the Duo Security service which is managed completely by Snowflake.
- B.** SSO is an authentication method that uses a username and password in the API request header.
- C.** SSO is an authentication method that allows a user to sign into multiple applications with a single set of credentials.
- D.** SSO is an authentication method which uses a pair of keys, a public key and a private key, to verify the identity of a user.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 217

What do these Query Profile metrics indicate about the query micro-partition pruning?

- A.** Since the bytes scanned is less than the bytes written to results some pruning has occurred.
- B.** No micro-partitions were pruned during execution.
- C.** The query used the local disk cache and pruning was not required.
- D.** The size of the data queried was too small to trigger micro-partition pruning.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 218

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 219

A directory table that references an external stage is being refreshed manually.

What will happen to any scheduled automatic-refresh operations occurring at the same time as the manual- refresh operations?

- A.** The automatic-refresh operations will continue on the established schedule.
- B.** The automatic-refresh operations will be performed more often.

- C. The automatic-refresh operations will be aborted.
- D. The manual-refresh operations and the automatic-refresh operations will both take place.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 220

Which default warehouse configuration has the highest precedence whenever a new session is created by a user?

- A. Default warehouse for the user
- B. Default warehouse specified on a CLI or in drivers/connectors parameters
- C. Default warehouse in the configuration file of the client utilities
- D. Default warehouse of the role assigned to the user

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 221

What does a Query Profile metric that shows excessive spillage indicate?

- A. Improper Common Table Expressions (CTEs)
- B. Poor query optimization
- C. Improper virtual warehouse sizing
- D. Expired temporary tables

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 222

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

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

Answer: ([SHOW ANSWER](#))

The UNDROP command in Snowflake can be used to directly restore Views and Tables. These objects, when dropped, are moved to a 'Recycle Bin' where they can be restored within a time limit before they are permanently deleted.

References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 223

A Snowflake user needs to optimize the definition of a secure view, but the user cannot see the view.

Which of the LEAST-PRIVILEGED access or role that should be granted to the user to complete this task?

- A. Grant the user the SYSADMIN role.
- B. Grant the user the SHOWFLAKE. object viewer database role.
- C. Grant the user the ownership privilege on the secure view.
- D. Grant the user the imported privileges privilege on the database.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 224

Which data types can be used in a Snowflake table that holds semi-structured data? (Select TWO).

- A. ARRAY
- B. BINARY
- C. TEXT
- D. VARIANT
- E. VARCHAK

Answer: ([SHOW ANSWER](#))

Snowflake supports semi-structured data types that can store complex data structures within a single column.

The **VARIANT** data type can hold structured and semi-structured data formats such as JSON, Avro, ORC, Parquet, and XML. The **ARRAY** data type is also used for semi-structured data and can store an ordered list of elements. These data types enable users to work with semi-structured data directly in Snowflake without needing to flatten the data into a relational schema first.

References: Snowflake Documentation on Semi-Structured Data Types

NEW QUESTION: 225

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 226

What metadata does Snowflake store concerning all rows stored in a micro-partition? (Select TWO).

- A. The number of distinct values for each column in the micro-partition
- B. The range of values for each of the rows in the micro-partition
- C. A count of the number of total values in the micro-partition
- D. The range of values for each of the columns in the micro-partition
- E. The range of values for each partition in the micro-partition

Answer: ([SHOW ANSWER](#))

answers have been corrected get the **newest** EduDump.com COF-C03 dumps with Test Engine here:

<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount

Code: **freecram**)

NEW QUESTION: 227

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

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

Answer: ([SHOW ANSWER](#))

Resource monitors in Snowflake can be configured at the account and virtual warehouse levels. They are used to track credit usage and control costs associated with running virtual warehouses. When certain thresholds are reached, resource monitors can trigger actions such as sending alerts or suspending warehouses to prevent excessive credit consumption. References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 228

When an ACCOUNTADMIN gives a user a custom role, what privilege or privileges is the user granted by default?

- A. Access to the PUBLIC role
- B. Access to all the objects owned by the USERADMIN role
- C. All privileges on objects allowed by the custom role
- D. All privileges that have been granted to the ACCOUNTADMIN

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 229

Which command is used to remove files from either external cloud storage or an internal staged

- A. DELETE
- B. REMOVE
- C. TRUNCATE
- D. DROP

Answer: ([SHOW ANSWER](#))

The REMOVE command in Snowflake is used to delete files from either external cloud storage locations or internal stages. This command helps manage staged files by removing them when they are no longer needed, which is useful for maintaining organization and managing storage costs.

Snowflake Documentation: REMOVE Command

NEW QUESTION: 230

Which Snowflake feature is used to track sensitive data to ensure compliance, identify data, and monitor resource usage?

- A. Data quality monitoring

- B. Multi-Factor Authentication (MFA)
- C. Object tagging
- D. Data classification

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 231

When cloning a database, what is cloned with the database? (Choose two.)

- A. Privileges on the database
- B. Existing child objects within the database
- C. Future child objects within the database
- D. Privileges on the schemas within the database
- E. Only schemas and tables within the database

Answer: ([SHOW ANSWER](#))

When cloning a database in Snowflake, the clone includes all privileges on the database as well as existing child objects within the database, such as schemas, tables, views, etc. However, it does not include future child objects or privileges on schemas within the database².

References = [COF-C03] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

NEW QUESTION: 232

A user is preparing to load data from an external stage

Which practice will provide the MOST efficient loading performance?

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

Answer: ([SHOW ANSWER](#))

Organizing files into logical paths can significantly improve the efficiency of data loading from an external stage. This practice helps in managing and locating files easily, which can be particularly beneficial when dealing with large datasets or complex directory structures¹.

NEW QUESTION: 233

Which access control entity in Snowflake can be created as part of a hierarchy within an account?

- A. Securable object
- B. Role
- C. Privilege
- D. User

Answer: ([SHOW ANSWER](#))

In Snowflake, a role is an access control entity that can be created as part of a hierarchy within an account. Roles are used to grant and manage privileges in a structured and scalable manner.

Understanding Roles:

Roles are logical entities that group privileges together.

They are used to control access to securable objects like tables, views, warehouses, and more.

Role Hierarchy:

Roles can be organized into a hierarchy, allowing for the inheritance of privileges.

A role higher in the hierarchy (parent role) can grant its privileges to a lower role (child role), simplifying privilege management.

Creating Roles:

Roles can be created using the CREATE ROLE command.

You can define parent-child relationships by granting one role to another.

Example Usage:

```
CREATE ROLE role1;
```

```
CREATE ROLE role2;
```

```
GRANT ROLE role1 TO role2;
```

In this example, role2 inherits the privileges of role1.

Benefits:

Simplifies privilege management: Hierarchies allow for efficient privilege assignment and inheritance.

Enhances security: Roles provide a clear structure for managing access control, ensuring that privileges are granted appropriately.

Snowflake Documentation: Access Control in Snowflake

Snowflake Documentation: Creating and Managing Roles

NEW QUESTION: 234

What is a key difference between the Snowflake CLI and SnowSQL?

A. SnowSQL is used for managing workloads and applications, while Snowflake CLI executes queries.

B. Snowflake CLI is primarily used for managing workloads and applications, while SnowSQL is used for executing operations.

C. SnowSQL supports only Data Definition Language, or DDL, operations, while Snowflake CLI supports only Data Manipulation Language, or DML, operations.

D. Snowflake CLI only supports query execution, while SnowSQL manages Python packages.

Answer: B (LEAVE A REPLY)

The correct answer is B. Snowflake CLI is primarily used for managing workloads and applications, while SnowSQL is used for executing operations .

SnowSQL is Snowflake's command-line client that is primarily used to connect to Snowflake and execute SQL statements. It is commonly used for SQL operations, including running queries, loading data, unloading data, and executing scripts.

Snowflake CLI is a broader command-line tool designed for managing Snowflake workloads and application development workflows. It supports tasks related to Snowflake Native Apps, Snowpark, Streamlit, stages, tasks, and other development or deployment activities.

Why B is correct:

Snowflake CLI is more focused on managing Snowflake workloads, projects, and applications. SnowSQL is primarily used as a SQL command-line client to execute SQL operations against Snowflake.

Why the other options are incorrect:

- A). This reverses the primary purposes of the tools. SnowSQL is not mainly for managing workloads and applications.
- C). SnowSQL is not limited to DDL, and Snowflake CLI is not limited to DML.
- D). Snowflake CLI is not limited to query execution, and SnowSQL is not a Python package management tool.

Official Snowflake documentation reference:

Snowflake documentation describes SnowSQL as the command-line client for connecting to Snowflake and executing SQL queries. Snowflake CLI is described as a command-line tool for developer workflows and managing Snowflake workloads and applications.

Reference: Snowflake Documentation - SnowSQL; Snowflake Documentation - Snowflake CLI; SnowPro Core Study Guide - Snowflake Account and Resource Management.

NEW QUESTION: 235

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: A,C (LEAVE A REPLY)

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: 236

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

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

Answer: (SHOW ANSWER)

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

NEW QUESTION: 237

In addition to the SELECT privilege, what other privilege must a role have to select from a table in database1. schema1?

- A. The USAGE privilege on the table.
- B. The OWNERSHIP privilege on the schema.
- C. The REFERENCES privilege on the table and the schema.
- D. The USAGE privilege on the database and the schema.

Answer: (SHOW ANSWER)

The correct answer is D. The USAGE privilege on the database and the schema .

To query a table in Snowflake, a role needs the SELECT privilege on the table. In addition, the role must have USAGE on the containing database and schema so it can resolve and access the object path.

Why D is correct:

For a table named:

database1.schema1.table1

A role generally needs:

USAGE ON DATABASE database1

USAGE ON SCHEMA database1.schema1

SELECT ON TABLE database1.schema1.table1

Without USAGE on the database and schema, the role cannot access the namespace that contains the table, even if SELECT exists on the table.

Why the other options are incorrect:

A). There is no USAGE privilege on a table for this purpose. SELECT is the table-level privilege needed to read data.

B). OWNERSHIP on the schema is not required to query a table.

C). REFERENCES is used for foreign key/reference-related behavior and is not required for simple SELECT queries.

Official Snowflake documentation reference:

Snowflake documentation explains that access to objects requires privileges on the object and appropriate USAGE privileges on parent containers, such as the database and schema.

Reference: Snowflake Documentation - Access control privileges; Snowflake Documentation - Overview of access control; SnowPro Core Study Guide - Security and Access Control.

NEW QUESTION: 238

The Information Schema and Account Usage Share provide storage information for which of the following objects? (Choose three.)

A. Users

B. Tables

C. Databases

D. Internal Stages

Answer: (SHOW ANSWER)

The Information Schema and Account Usage Share in Snowflake provide metadata and historical usage data for various objects within a Snowflake account. Specifically, they offer storage information for Tables, Databases, and Internal Stages. These schemas contain views and table functions that allow users to query object metadata and usage metrics, such as the amount of data stored and historical activity.

Tables: The storage information includes data on the daily average amount of data in database tables.

Databases: For databases, the storage usage is calculated based on all the data contained within the database, including tables and stages.

Internal Stages: Internal stages are locations within Snowflake for temporarily storing data, and their storage usage is also tracked.

References: The information is verified according to the SnowPro Core Certification Study Guide and Snowflake documentation

NEW QUESTION: 239

When resizing a currently running virtual warehouse, what happens to in-progress queries?

- A. In-progress queries will be cancelled.
- B. Queries will continue running, using the new warehouse size.
- C. Queries will continue running, using the current warehouse size.
- D. Queries will be paused and automatically reprocessed using the new warehouse size.

Answer: (SHOW ANSWER)

The correct answer is C. Queries will continue running, using the current warehouse size .

When a virtual warehouse is resized while it is running, the resize does not affect queries that are already in progress.

Running queries continue using the warehouse resources that were available when they started. The new warehouse size applies to queries that start after the resize operation.

Why C is correct:

Snowflake does not cancel or restart existing queries when a warehouse is resized. In-progress queries continue executing with the current warehouse size.

Why the other options are incorrect:

- A). Resizing a warehouse does not automatically cancel in-progress queries.
- B). In-progress queries do not switch to the new warehouse size.
- D). Queries are not paused and reprocessed after resizing.

Official Snowflake documentation reference:

Snowflake documentation explains that resizing a running warehouse affects only new queries and queued queries.

Currently running queries continue to use the warehouse size that was in effect when they started.

Reference: Snowflake Documentation - Resizing a warehouse; Snowflake Documentation - Virtual warehouses; SnowPro Core Study Guide - Snowflake Account and Resource Management.

NEW QUESTION: 240

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: B,D (LEAVE A REPLY)

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-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 241

Which will enable temporary scaling for queries that contain complex computations that may consume many credits?

- A. Search Optimization Service
- B. Virtual warehouse scaling policy

- C. Query Acceleration Service
- D. Multi-cluster virtual warehouse

Answer: (SHOW ANSWER)

The correct answer is C. Query Acceleration Service .

The Query Acceleration Service allows Snowflake to use additional compute resources to accelerate eligible queries. It is especially useful for queries that scan large amounts of data or include expensive operations that can benefit from extra parallel processing.

Why C is correct:

Query Acceleration Service temporarily uses additional Snowflake-managed compute resources for eligible queries. This can improve performance without permanently resizing the virtual warehouse.

Why the other options are incorrect:

- A). Search Optimization Service improves highly selective lookup queries by maintaining search access paths, but it does not temporarily scale compute for expensive queries.
- B). A warehouse scaling policy controls how multi-cluster warehouses add or remove clusters, but it is not the feature designed to accelerate individual expensive queries.
- D). A multi-cluster warehouse improves concurrency by adding clusters, but it does not split or accelerate a single complex query in the way Query Acceleration Service can.

Official Snowflake documentation reference:

Snowflake documentation describes Query Acceleration Service as a feature that can improve query performance by using additional compute resources for eligible parts of queries.

Reference: Snowflake Documentation - Query Acceleration Service; Snowflake Documentation - Warehouses; SnowPro Core Study Guide - Performance Optimization.

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:

<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF Special Discount**

Code: freecram)

NEW QUESTION: 242

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: (SHOW ANSWER)

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: 243

What is the LEAST-PRIVILEGED role that can create a share in Snowflake by default?

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 244

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: ([SHOW ANSWER](#))

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

NEW QUESTION: 245

Which types of subqueries does Snowflake support? (Select TWO).

- A. Uncorrelated scalar subqueries in WHERE clauses
- B. Uncorrelated scalar subqueries in any place that a value expression can be used
- C. EXISTS, ANY / ALL, and IN subqueries in WHERE clauses: these subqueries can be uncorrelated only
- D. EXISTS, ANY / ALL, and IN subqueries in where clauses: these subqueries can be correlated only
- E. EXISTS, ANY /ALL, and IN subqueries in WHERE clauses: these subqueries can be correlated or uncorrelated

Answer: ([SHOW ANSWER](#))

Snowflake supports a variety of subquery types, including both correlated and uncorrelated subqueries. The correct answers are B and E, which highlight Snowflake 's flexibility in handling subqueries within SQL queries.

Uncorrelated Scalar Subqueries: These are subqueries that can execute independently of the outer query. They return a single value and can be used anywhere a value expression is allowed, offering great flexibility in SQL queries.

EXISTS, ANY/ALL, and IN Subqueries: These subqueries are used in WHERE clauses to filter the results of the main query based on the presence or absence of matching rows in a subquery. Snowflake supports both correlated and uncorrelated versions of these subqueries, providing powerful tools for complex data analysis scenarios.

Examples and Usage:

Uncorrelated Scalar Subquery:

```
SELECT * FROM Employees WHERE salary > (SELECT AVG(salary) FROM Employees);
```

Correlated EXISTS Subquery:

```
SELECT * FROM Orders o WHERE EXISTS (SELECT 1 FROM customer c WHERE c.id = o.customer_id AND c.
```

region = ' North America ');

Reference: For comprehensive details on the types of subqueries supported by Snowflake and examples of their usage, consult the Snowflake documentation on subqueries: <https://docs.snowflake.com/en/sql-reference/constructs/subqueries.html>

NEW QUESTION: 246

By default, which Snowflake role has the ability to manage grants?

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

Answer: (SHOW ANSWER)

The correct answer is A. SECURITYADMIN .

In Snowflake's default role hierarchy, the SECURITYADMIN role is responsible for managing security- related objects and grants. This role is granted the global MANAGE GRANTS privilege, which allows it to manage object grants across the account.

Why A is correct:

SECURITYADMIN is the built-in role designed to manage users, roles, and grants. By default, it has the ability to grant and revoke privileges.

Why the other options are incorrect:

- B). USERADMIN is used primarily for user and role management, but it does not have the same broad default grant management authority as SECURITYADMIN.
- C). SYSADMIN is typically used to create and manage account objects such as warehouses, databases, schemas, and tables, but it is not the default security grant administrator.
- D). PUBLIC is automatically granted to every user and should not be used for administrative grant management.

Official Snowflake documentation reference:

Snowflake documentation describes SECURITYADMIN as the role that manages grants globally and inherits the privileges of USERADMIN. It includes the MANAGE GRANTS privilege by default.

Reference: Snowflake Documentation - System-defined roles; Snowflake Documentation - Access control overview; SnowPro Core Study Guide - Security and Access Control.

NEW QUESTION: 247

What calculation of metrics found in the QUERY_HISTORY view can be used to determine whether optimizing virtual warehouse caching will improve the performance of a long-running query?

- A. Multiply the partitions scanned and the percentage scanned from cache.
- B. Multiply the bytes scanned and the percentage scanned from cache.
- C. Multiply the bytes written and the percentage scanned from cache.
- D. Multiply the bytes scanned and the percentage scanned from cache miss.

Answer: (SHOW ANSWER)

The useful calculation is based on BYTES_SCANNED and PERCENTAGE_SCANNED_FROM_CACHE.

These metrics help estimate how much scanned data came from the warehouse cache. A low cache percentage on a query that scans many bytes may indicate that warehouse cache optimization could help.

Corrected note: the original options had OCR errors. The best matching valid answer is B .

Reference: Snowflake Documentation - QUERY_HISTORY view; Warehouse cache; Query performance monitoring.

NEW QUESTION: 248

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 2 < 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: (SHOW ANSWER)

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: 249

A virtual warehouse is created using the following command:

```
Create warehouse my_WH with  
warehouse_size = MEDIUM  
min_cluster_count = 1  
max_cluster_count = 1  
auto_suspend = 60  
auto_resume = true;
```

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



What action should be taken to address this situation?

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

Answer: ([SHOW ANSWER](#))

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

References = Snowflake Documentation on Multi-cluster Warehouses¹

NEW QUESTION: 250

What MINIMUM permissions are required to create a pipe in Snowflake? (Select TWO).

- A. OWNERSHIP on the table in the pipe definition
- B. CREATE PIPE at the schema level
- C. SELECT and INSERT on the table in the pipe definition
- D. CREATE SCHEMA at the database level

E. USAGE on the virtual warehouse

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 251

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: ([SHOW ANSWER](#))

[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: 252

Which functions can be used to extract values from a VARIANT column by specifying a hierarchy name?

Select TWO.

A. GET()

B. TO_JSON()

C. GET_PATH()

D. PARSE_JSON()

E. LATERAL FLATTEN

Answer: ([SHOW ANSWER](#))

The correct answers are A. GET() and C. GET_PATH() .

Snowflake supports functions for extracting values from semi-structured data stored in a VARIANT, OBJECT, or ARRAY column. GET() and GET_PATH() can retrieve values by specifying keys, indexes, or paths.

Why A is correct:

GET() extracts a value from an object or array. For an object, it can retrieve the value associated with a specified key.

Example:

```
SELECT GET(src, ' customer ' )  
FROM my_table;
```

Why C is correct:

GET_PATH() extracts a value from semi-structured data using a path expression. This is useful when specifying a hierarchy or nested path.

Example:

```
SELECT GET_PATH(src, ' customer.name ' )  
FROM my_table;
```

Why the other options are incorrect:

B). TO_JSON() converts a VARIANT value to a JSON string. It does not extract values by hierarchy name.

D). `PARSE_JSON()` parses a string into a `VARIANT`. It does not extract an existing nested value.

E). `LATERAL FLATTEN` expands arrays or objects into rows. It is useful for traversing semi-structured data, but it is not a scalar extraction function that retrieves a value by specifying a hierarchy name.

Official Snowflake documentation reference:

Snowflake documentation describes `GET()` as a function that extracts a value from an object or array, and

`GET_PATH()` as a function that extracts a value from semi-structured data using a path name.

Reference: Snowflake Documentation - `GET`; Snowflake Documentation - `GET_PATH`; Snowflake Documentation - Querying semi-structured data; SnowPro Core Study Guide - Working with Semi- Structured Data.

NEW QUESTION: 253

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 254

Which privilege is needed for a Snowflake user to see the definition of a secure view?

- A. `OWNERSHIP`
- B. `MODIFY`
- C. `CREATE`
- D. `USAGE`

Answer: ([SHOW ANSWER](#))

To see the definition of a secure view in Snowflake, the minimum privilege required is `OWNERSHIP` of the view.

Ownership grants the ability to view the definition as well as to modify or drop the view. Secure views are designed to protect sensitive data, and thus the definition of these views is restricted to users with sufficient privileges to ensure data security.

References:

Snowflake Documentation: Secure Views

NEW QUESTION: 255

Use of which feature or setting consumes Snowflake credits, but generates no storage costs?

- A. Materialized view
- B. Query acceleration service
- C. Table clustering
- D. Search optimization service

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 256

Which object type is granted permissions for reading a table?

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

Answer: ([SHOW ANSWER](#))

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>

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here: <https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (**1187** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 257

What are valid sub-clauses to the OVER clause for a window function? (Select TWO).

- A. GROUP BY
- B. LIMIT
- C. ORDER BY
- D. PARTITION BY
- E. UNION ALL

Answer: C,D ([LEAVE A REPLY](#))

Valid sub-clauses to the OVER clause for a window function in SQL are:

C). ORDER BY: This clause specifies the order in which the rows in a partition are processed by the window function. It is essential for functions that depend on the row order, such as ranking functions.

D). PARTITION BY: This clause divides the result set into partitions to which the window function is applied.

Each partition is processed independently of other partitions, making it crucial for functions that compute values across sets of rows that share common characteristics.

These clauses are fundamental to defining the scope and order of data over which the window function operates, enabling complex analytical computations within SQL queries.

References:

Snowflake Documentation: Window Functions

NEW QUESTION: 258

How should a Snowflake user access a third-party SaaS service to process unstructured data?

- A. Use internal functions.
- B. Use external functions.
- C. Use an API gateway.
- D. Use process functions.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 259

Which service or feature in Snowflake is used to improve the performance of certain types of lookup and analytical queries that use an extensive set of WHERE conditions?

- A. Data classification
- B. Query acceleration service
- C. Search optimization service
- D. Tagging

Answer: ([SHOW ANSWER](#))

The Search Optimization Service in Snowflake is designed to improve the performance of specific types of queries, particularly those involving extensive sets of WHERE conditions. By maintaining a search index on tables, this service can accelerate lookup and analytical queries, making it a valuable feature for optimizing query performance and reducing execution times for complex searches.

References:

Snowflake Documentation: Search Optimization Service

NEW QUESTION: 260

What is a key characteristic of managed access schemas?

- A. Object owners can grant privileges on their objects to other roles.
- B. Any user with object ownership can manage future grants independently.
- C. Only the schema owner or roles with the MANAGE GRANTS privilege can grant privileges on objects.
- D. The global MANAGE GRANTS privilege is not applicable for granting additional privileges to other roles.

Answer: C ([LEAVE A REPLY](#))

The correct answer is C. Only the schema owner or roles with the MANAGE GRANTS privilege can grant privileges on objects .

In a managed access schema , privilege management is centralized. Individual object owners do not manage grants on their own objects. Instead, grants on objects inside the schema are managed by the schema owner or a role with the global MANAGE GRANTS privilege.

Why C is correct:

Managed access schemas are designed to simplify and centralize access control. Even if a role owns a table, view, or other object inside the managed access schema, that role cannot independently grant privileges on that object unless it is also the schema owner or has **MANAGE GRANTS**.

Why the other options are incorrect:

- A). This describes behavior in regular schemas, not managed access schemas.
- B). Object ownership alone is not enough to manage future grants independently in a managed access schema.
- D). The global **MANAGE GRANTS** privilege is applicable and allows privilege management.

Official Snowflake documentation reference:

Snowflake documentation explains that in managed access schemas, object owners lose the ability to make grant decisions. Only the schema owner or a role with the **MANAGE GRANTS** privilege can grant privileges on objects in the schema.

Reference: Snowflake Documentation - Managed access schemas; Snowflake Documentation - Access control privileges; SnowPro Core Study Guide - Security and Access Control.

NEW QUESTION: 261

Which function unloads data from a relational table to JSON?

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

Answer: (SHOW ANSWER)

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-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 262

What notation methods are supported in Snowflake to access elements in a JSON object? Select TWO.

- A. Classless Inter-Domain Routing, or CIDR, notation
- B. Vertical bar notation
- C. Bracket notation
- D. Slash notation
- E. Dot notation

Answer: (SHOW ANSWER)

The correct answers are C. Bracket notation and E. Dot notation .

Snowflake supports querying semi-structured data stored in columns such as **VARIANT**, **OBJECT**, and **ARRAY**. When working with JSON objects, users can access elements using either dot notation or bracket notation .

Why C is correct:

Bracket notation is used to access JSON object elements by enclosing the key name in brackets.

Example:

```
SELECT src[ ' customer ' ][ ' name ' ]  
FROM json_table;
```

Bracket notation is especially useful when JSON keys contain special characters, spaces, or case-sensitive names.

Why E is correct:

Dot notation is used to access JSON object elements using a period between object levels.

Example:

```
SELECT src:customer.name  
FROM json_table;
```

This notation is commonly used for simple JSON key paths.

Why the other options are incorrect:

- A). CIDR notation is used for IP address ranges, not JSON element access.
- B). Vertical bar notation is not a Snowflake method for accessing JSON object elements.
- D). Slash notation is not a Snowflake method for accessing JSON object elements.

Official Snowflake documentation reference:

Snowflake documentation explains that semi-structured data elements in VARIANT columns can be accessed using dot notation and bracket notation.

Reference: Snowflake Documentation - Querying semi-structured data; Snowflake Documentation - Dot notation and bracket notation; SnowPro Core Study Guide - Working with Semi-Structured Data.

NEW QUESTION: 263

When a column in a table is dropped, what is the impact on the table 's micro-partitions?

- A. New micro-partitions will be created that include all remaining columns and data.
- B. The table 's micro-partitions will be updated to reflect the removal of the data in the dropped column.
- C. A new table will be created, without the data in the dropped column, and new micro-partitions will be created.
- D. The existing micro-partitions will not be changed or impacted.

Answer: (SHOW ANSWER)

The correct answer is D. The existing micro-partitions will not be changed or impacted .

Important correction:

The provided answer B is not correct.

Snowflake stores table data in immutable micro-partitions. Because micro-partitions are immutable, Snowflake does not physically rewrite all existing micro-partitions immediately when a column is dropped.

Instead, metadata is updated so that the dropped column is no longer visible or accessible in the table definition.

Why D is correct:

Dropping a column is a metadata operation. The underlying existing micro-partitions are not immediately rewritten to remove the column data. Snowflake's immutable micro-partition design avoids expensive full- table rewrites for this type of operation.

Why the other options are incorrect:

- A). Snowflake does not immediately create new micro-partitions containing only the remaining columns.
- B). Existing micro-partitions are immutable and are not updated in place.
- C). Snowflake does not automatically create a new table and rewrite all remaining data when a column is dropped.

Official Snowflake documentation reference:

Snowflake documentation describes micro-partitions as immutable storage units. Table metadata operations, such as dropping columns, do not immediately rewrite existing micro-partitions.

Reference: Snowflake Documentation - Micro-partitions and data clustering; Snowflake Documentation - ALTER TABLE ... DROP COLUMN; SnowPro Core Study Guide - Snowflake Architecture.

NEW QUESTION: 264

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

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

Answer: ([SHOW ANSWER](#))

Snowflake's Time Travel feature is used for both querying historical data in tables and restoring and cloning historical data in databases, schemas, and tables³. It allows users to access historical data within a defined period (1 day by default, up to 90 days for Snowflake Enterprise Edition) and is a key feature for data recovery and management.

References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 265

What is the PRIMARY factor that determines the cost of using a virtual warehouse in Snowflake?

- A. The type of SQL statements executed
- B. The number of tables or databases queried
- C. The amount of data stored in the warehouse
- D. The length of time the compute resources in each cluster run

Answer: ([SHOW ANSWER](#))

The primary factor that determines the cost of using a virtual warehouse in Snowflake is D. The length of time the compute resources in each cluster run. Snowflake's pricing model for compute usage is based on the concept of Snowflake credits, which are consumed based on the time virtual warehouses are running and the size of the warehouses. The more compute resources are utilized and the longer they run, the more credits are consumed, which directly impacts the cost.

References:

Snowflake Documentation on Virtual Warehouses: Understanding and Managing Warehouse Credit Usage

NEW QUESTION: 266

What privileges or user roles are required to get or request access to data that is available in the Snowflake Marketplace? (Select TWO).

- A. ACCOUNTADMIN
- B. MANAGE SHARE
- C. SECURITYADMIN
- D. IMPORT SHARE
- E. USERADMIN

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 267

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 268

Which Snowflake security feature has the HIGHEST precedence in the security feature hierarchy?

- A. Network policies
- B. Snowflake OAuth
- C. Account roles
- D. Masking policies

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 269

When a Snowflake user loads CSV data from a stage, which copy into < table > command guideline should they follow?

- A. The CSV field delimiter must be a comma character (,),
- B. The number of columns in each row should be consistent.
- C. The data file in the stage must be in a compressed format.
- D. The data file must have the same number of columns as the target table.

Answer: ([SHOW ANSWER](#))

When using the COPY INTO < table > command to load CSV data from a stage into a Snowflake table, one of the crucial guidelines to follow is that the number of columns in each row within the CSV file should be consistent. This ensures data integrity and allows for a smooth data loading process, as Snowflake expects each row in the CSV file to map directly to a row in the target table based on the number of columns. If there ' s a mismatch in the number of columns between any row in the file and the target table, Snowflake might return an error or produce unexpected results during the load operation.

It ' s important to note that while the CSV field delimiter can be specified to something other than a comma, ensuring the consistency in the number of columns across all rows is fundamental to successfully loading data.

References:

Snowflake Documentation on Loading Data: Loading CSV Data

NEW QUESTION: 270

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: ([SHOW ANSWER](#))

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: 271

Which activities are included in the Cloud Services layer? {Select TWO}.

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

Answer: ([SHOW ANSWER](#))

The Cloud Services layer in Snowflake is responsible for a wide range of services that facilitate the management and use of Snowflake, including:

D). User authentication: This service handles identity and access management, ensuring that only authorized users can access Snowflake resources.

E). Infrastructure management: This service manages the allocation and scaling of resources to meet user demands, including the management of virtual warehouses, storage, and the orchestration of query execution.

These services are part of Snowflake 's fully managed, cloud-based architecture, which abstracts and automates many of the complexities associated with data warehousing.

References:

Snowflake Documentation: Overview of Snowflake Cloud Services

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:

<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount

Code: **freecram**)

NEW QUESTION: 272

Which roles are needed to access Document AI? Select TWO.

- A. PUBLIC
- B. SYSADMIN
- C. ACCOUNTADMIN

D. OBJECT_VIEWER

E. DOCUMENT_INTELLIGENCE_CREATOR

Answer: ([SHOW ANSWER](#))

The correct answers are C. ACCOUNTADMIN and E. DOCUMENT_INTELLIGENCE_CREATOR .

Document AI is part of Snowflake Cortex and uses specific privileges/roles for access and administration. The ACCOUNTADMIN role is commonly required for initial setup or granting required privileges, while DOCUMENT_INTELLIGENCE_CREATOR is used to create and work with Document AI objects.

Why C is correct:

ACCOUNTADMIN has the account-level authority needed to enable or manage required access for Document AI features.

Why E is correct:

DOCUMENT_INTELLIGENCE_CREATOR is the Snowflake role associated with creating Document AI- related objects and using Document AI capabilities.

Why the other options are incorrect:

A). PUBLIC is available to all users but does not provide Document AI creation privileges.

B). SYSADMIN manages many account objects, but it is not the specific Document AI role in the answer set.

D). OBJECT_VIEWER is a Snowflake database role for object metadata visibility, not Document AI access.

Official Snowflake documentation reference:

Snowflake documentation for Document AI describes the required access model and includes the DOCUMENT_INTELLIGENCE_CREATOR role for creating Document AI objects, with account administrator involvement for setup.

Reference: Snowflake Documentation - Document AI; Snowflake Cortex; SnowPro Core Study Guide - Snowflake Architecture and Security.

NEW QUESTION: 273

What type of query benefits the MOST from search optimization?

A. A query that uses only disjunction (i.e., OR) predicates

B. A query that includes analytical expressions

C. A query that uses equality predicates or predicates that use IN

D. A query that filters on semi-structured data types

Answer: ([SHOW ANSWER](#))

Search optimization in Snowflake is designed to improve the performance of queries that are selective and involve point lookup operations using equality and IN predicates. It is particularly beneficial for queries that access columns with a high number of distinct values¹.

References = [COF-C03] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

NEW QUESTION: 274

Which type of role is designed to control access to objects that belong to the same logical group of tables, views, and schemas?

A. Account

B. Database

- C. Service
- D. Application

Answer: ([SHOW ANSWER](#))

The correct answer is B. Database .

A database role is designed to manage access to objects within a database. Database roles are useful for grouping privileges on objects that belong to the same logical database-level domain, such as related schemas, tables, and views.

Why B is correct:

Database roles are created inside a database and can be granted privileges on objects in that database. They help organize access control for database-specific objects.

Example concept:

```
CREATE DATABASE ROLE sales_db.sales_reader;  
GRANT USAGE ON SCHEMA sales_db.reporting TO DATABASE ROLE sales_db.sales_reader;  
GRANT SELECT ON ALL TABLES IN SCHEMA sales_db.reporting TO DATABASE ROLE sales_db.  
sales_reader;
```

The database role can then be granted to an account role.

Why the other options are incorrect:

A). Account roles are account-level roles and can hold privileges across many databases and account objects.

They are broader than the logical grouping described in the question.

C). Service roles are associated with service-related access, not general database object access control.

D). Application roles are used by Snowflake Native Apps to manage access inside an application, not to generally control access to a logical group of database objects.

Official Snowflake documentation reference:

Snowflake documentation describes database roles as roles scoped to a database and used to manage privileges on database objects.

Reference: Snowflake Documentation - Database roles; Snowflake Documentation - Access control overview; SnowPro Core Study Guide - Security and Access Control.

NEW QUESTION: 275

What is required to set a stored procedure to EXECUTE AS? (Select TWO).

- A. SCHEMA_USER
- B. USER
- C. Owner ' s rights
- D. ACCOUNTADMIN role
- E. Caller ' s rights

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 276

What is the primary purpose of using a masking policy in Snowflake?

- A. To protect sensitive data from unauthorized access when queries are run.
- B. To automatically encrypt sensitive data when data is stored in Snowflake.

- C. To protect multiple columns that have different data types in a given table.
- D. To protect both column-level and row-level data.

Answer: ([SHOW ANSWER](#))

Masking policies in Snowflake are designed to protect sensitive information by dynamically hiding or obfuscating data based on the role of the user executing the query. This helps enforce data privacy and security by allowing only authorized users to see sensitive information. Masking policies do not encrypt data but apply rules to limit data visibility, ensuring sensitive data is protected during query execution without altering the underlying data.

NEW QUESTION: 277

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 ([LEAVE A REPLY](#))

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: 278

What is a benefit of using system tags?

- A. To manage compute resources
- B. To support data governance
- C. To increase query execution speed
- D. To control data backup frequency

Answer: ([SHOW ANSWER](#))

System tags help enable data classification for governance, privacy, and compliance. For example, Snowflake uses system-defined tags like SNOWFLAKE.CORE.DATA_CLASSIFICATION to detect and label sensitive data (e.g., PII, financial info).

Reference:

Snowflake Docs: System Tags

NEW QUESTION: 279

Which commands extract the value for a key, key_1, from an object column? Select TWO.

- A. SELECT column:key_1 FROM my_table;
- B. SELECT column.key_1 FROM my_table;
- C. SELECT column[' key_1 '] FROM my_table;
- D. SELECT column(' key_1 ') FROM my_table;

E. `SELECT key_1 FROM column::string AS key_1;`

Answer: A,C (LEAVE A REPLY)

The correct answers are A. `SELECT column:key_1 FROM my_table;` and C. `SELECT column[ ' key_1 ' ] FROM my_table;` .

Snowflake supports path notation for querying semi-structured data stored in columns such as VARIANT, OBJECT, and ARRAY. For a JSON object key, Snowflake supports colon path notation and bracket notation.

Why A is correct:

The colon path notation accesses a field inside a semi-structured column.

Example:

```
SELECT column:key_1
FROM my_table;
```

This extracts the value associated with the key `key_1`.

Why C is correct:

Bracket notation also accesses an object element by key.

Example:

```
SELECT column[ ' key_1 ' ]
FROM my_table;
```

This is especially useful when key names require quoting or contain special characters.

Why the other options are incorrect:

B). Dot notation is supported for nested traversal after the initial colon path, such as `column:key_1.subkey`, but `column.key_1` by itself is interpreted as regular SQL object/column qualification, not as extracting a key from a semi-structured column.

D). Parentheses are not valid JSON object key extraction syntax in Snowflake.

E). Casting a column to string and aliasing it does not extract the value for the key `key_1`.

Official Snowflake documentation reference:

Snowflake documentation explains that elements in semi-structured data can be accessed using colon path notation and bracket notation. Dot notation may also be used for nested paths after the semi-structured column reference.

Reference: Snowflake Documentation - Querying semi-structured data; Snowflake Documentation - Traversing semi-structured data; SnowPro Core Study Guide - Working with Semi-Structured Data.

NEW QUESTION: 280

How does the search optimization service improve query performance?

A. By clustering the tables

B. By creating a persistent data structure

C. By using caching

D. By optimizing the use of micro-partitions

Answer: (SHOW ANSWER)

The Search Optimization Service in Snowflake enhances query performance by creating a persistent data structure that enables faster access to specific data, particularly for queries with selective filters on columns not often used in clustering. This persistent structure accelerates data retrieval without depending on clustering or caching, thereby improving response times for targeted queries.

Snowflake 's micro-partitioning automatically manages table structure, but search optimization allows further enhancement for certain high-frequency, specific access patterns.

NEW QUESTION: 281

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: ([SHOW ANSWER](#))

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: 282

What Multi-Factor Authentication, or MFA, service is provided by Snowflake?

- A. Microsoft Entra ID
- B. Google Authenticator
- C. Duo Security
- D. Microsoft Authenticator

Answer: ([SHOW ANSWER](#))

The correct answer is C. Duo Security .

Snowflake provides integrated Multi-Factor Authentication through Duo Security . MFA adds an additional verification step beyond username and password authentication, improving account security.

Why C is correct:

Snowflake's native MFA integration uses Duo Security. Users can enroll in MFA and authenticate using Duo- supported methods.

Why the other options are incorrect:

- A). Microsoft Entra ID can be used as an identity provider for federated authentication, but it is not the MFA service provided natively by Snowflake.
- B). Google Authenticator is a common authenticator app, but it is not Snowflake's native MFA service.
- D). Microsoft Authenticator is also a common MFA app, but Snowflake's built-in MFA service is Duo Security.

Official Snowflake documentation reference:

Snowflake documentation describes Snowflake MFA as being powered by Duo Security.

Reference: Snowflake Documentation - Multi-factor authentication; Snowflake Documentation - User authentication; SnowPro Core Study Guide - Security and Access Control.

NEW QUESTION: 283

If queries start to queue in a multi-cluster virtual warehouse, an additional compute cluster starts immediately under what setting?

- A. Auto-scale mode

- B. Maximized mode
- C. Economy scaling policy
- D. Standard scaling policy

Answer: ([SHOW ANSWER](#))

In Snowflake, when queries begin to queue in a multi-cluster virtual warehouse, an additional compute cluster starts immediately if the warehouse is set to auto-scale mode. This mode allows Snowflake to automatically add or resume additional clusters as soon as the workload increases, and similarly, shut down or pause the additional clusters when the load decreases

NEW QUESTION: 284

Which COPY INTO < location > command options impact the number of files generated? Select TWO.

- A. SINGLE
- B. RECORD_DELIMITER
- C. MAX_FILE_SIZE
- D. INCLUDE_QUERY_ID
- E. DETAILED_OUTPUT

Answer: ([SHOW ANSWER](#))

The correct answers are A. SINGLE and C. MAX_FILE_SIZE .

When unloading data from Snowflake using COPY INTO < location > , Snowflake can create one or more output files. Some options influence how many files are generated.

Why A is correct:

The SINGLE option specifies whether Snowflake generates a single output file or multiple output files. When SINGLE = TRUE, Snowflake unloads data into a single file, subject to command and file size behavior.

Why C is correct:

The MAX_FILE_SIZE option specifies the maximum size of each file generated during unloading. Smaller maximum file sizes can result in more files, while larger maximum file sizes can result in fewer files.

Why the other options are incorrect:

B). RECORD_DELIMITER controls how records are separated inside the output file. It does not control how many files are created.

D). INCLUDE_QUERY_ID controls whether the query ID is included in unloaded file names. It affects naming behavior, not the number of files generated.

E). DETAILED_OUTPUT controls the level of output information returned by the command. It does not control the number of unloaded files.

Official Snowflake documentation reference:

Snowflake documentation for COPY INTO < location > describes SINGLE as controlling whether a single file is generated and MAX_FILE_SIZE as controlling the maximum size of unloaded files.

Reference: Snowflake Documentation - COPY INTO < location > ; Snowflake Documentation - Data unloading considerations; SnowPro Core Study Guide - Data Loading and Unloading.

NEW QUESTION: 285

What is the first step when creating a Data Exchange in Snowflake?

- A. Create an organization ID.
- B. Run the CREATE LISTING command.
- C. Run the CREATE SHARE command.
- D. Contact Snowflake Support.

Answer: ([SHOW ANSWER](#))

The correct answer is D. Contact Snowflake Support .

Important correction:

The provided answer A. Create an organization ID is not the best answer. A Snowflake organization is required for certain marketplace and sharing features, but users do not typically "create an organization ID" as the first procedural step for setting up a Data Exchange.

Why D is correct:

A Data Exchange is a private data hub that allows providers and consumers to securely share data. Creating or enabling a Data Exchange typically requires Snowflake involvement. The first step is to contact Snowflake Support or a Snowflake representative to enable or configure the Data Exchange.

Why the other options are incorrect:

- A). An organization identifier may be involved in account and organization-level administration, but creating an organization ID is not the first step to create a Data Exchange.
- B). CREATE LISTING is used for creating listings after the provider/share setup exists.
- C). CREATE SHARE creates a share object, but it is not the first step for creating the Data Exchange itself.

Official Snowflake documentation reference:

Snowflake documentation describes Data Exchange as a controlled data sharing hub and indicates that setting up a Data Exchange requires working with Snowflake.

Reference: Snowflake Documentation - Data Exchange; Snowflake Documentation - Secure Data Sharing; SnowPro Core Study Guide - Data Protection and Governance.

NEW QUESTION: 286

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: ([SHOW ANSWER](#))

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

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here: <https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 287

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: (SHOW ANSWER)

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-C03] SnowPro Core Certification Exam Study Guide

Snowflake Documentation12

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

NEW QUESTION: 288

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: (SHOW ANSWER)

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: 289

When unloading data to an external stage, what is the MAXIMUM file size supported?

- A. 1 GB
- B. 5 GB
- C. 10 GB
- D. 16 GB

Answer: ([SHOW ANSWER](#))

When unloading data to an external stage, the maximum file size supported is 5 GB. This limit ensures efficient data transfer and management within Snowflake's architecture

NEW QUESTION: 290

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 ([LEAVE A REPLY](#))

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-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 291

Why is a federated environment used for user authentication in Snowflake?

- A. To enhance data security and privacy
- B. To provide real-time monitoring of user activities
- C. To separate user authentication from user access
- D. To enable direct integration with external databases

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 292

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

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

Answer: ([SHOW ANSWER](#))

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: 293

Which role can create and manage Snowflake accounts?

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 294

A Snowflake user needs to share unstructured data from an internal stage to a reporting tool that does not have Snowflake access.

Which file function should be used?

- A. BUILD_SCOPED_FILE_URL
- B. BUILD_STAGE_FILE_URL
- C. GET_PREIGNED_URL
- D. GET STAGE LOCATION

Answer: ([SHOW ANSWER](#))

The GET_PREIGNED_URL function in Snowflake generates a presigned URL for a file stored in an internal stage, allowing direct access to the file without requiring Snowflake access. This feature is particularly useful for sharing unstructured data with external applications or tools that do not have direct access to Snowflake. The presigned URL provides temporary access to the file, making it an ideal solution for securely sharing unstructured data from an internal stage with a reporting tool or any other external application.

References:

Snowflake Documentation: Generating Presigned URLs

NEW QUESTION: 295

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

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

Answer: B ([LEAVE A REPLY](#))

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

References:

Snowflake Documentation on Query Results Cache

SnowPro Core Certification Study Guide

NEW QUESTION: 296

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: ([SHOW ANSWER](#))

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 query.

NEW QUESTION: 297

A JSON document is stored in the source_column of type VARIANT. The document has an array called elements. The array contains the name key that has a string value. How can a Snowflake user extract the name from the first element?

- A. Source_column.element[1]:name
- B. Source_column.element[0]:name
- C. Source_column.element[1].name
- D. Source_column.element[0].name

Answer: ([SHOW ANSWER](#))

In Snowflake, when dealing with semi-structured data such as a JSON document stored in a VARIANT column, the proper syntax to extract a value is to use the column name followed by the path to the specific element. Since arrays in JSON are zero-indexed, the first element is referenced with [0]. Therefore, to extract the name from the first element of the elements array, the correct syntax is Source_column:elements[0].name.

References: Snowflake Documentation on Semi-Structured Data

NEW QUESTION: 298

What will be the result of this command?

```
SELECT value FROM LATERAL FLATTEN(input => parse_json(' [1, 2, 3] '));
```

- A. It will return a single row with the values 1, 2, and 3 concatenated.
- B. It will return three rows with the values 1, 2, and 3.
- C. It will return a single row with three columns with values 1, 2, and 3.
- D. It will return a single row with the array [1, 2, 3].

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 299

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 ([LEAVE A REPLY](#))

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: 300

What is a directory table in Snowflake?

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 301

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

- A. OBJECT_CONSTRUCT
- B. PARSE_JSON
- C. OBJECT_AGG
- D. FLATTEN

Answer: ([SHOW ANSWER](#))

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and

answers have been corrected get the **newest** EduDump.com COF-C03 dumps with Test Engine here:

<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount

Code: **freecram**)

NEW QUESTION: 302

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

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

Answer: ([SHOW ANSWER](#))

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

References:

Snowflake Documentation on Pipes

SnowPro Core Certification Study Guide

<https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro.html>

NEW QUESTION: 303

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: ([SHOW ANSWER](#))

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 use.

NEW QUESTION: 304

To switch between a Gen2 standard warehouse and a Snowpark-optimized warehouse, which warehouse object properties must be changed? Select TWO.

- A. RESOURCE_MONITOR
- B. WAREHOUSE_TYPE
- C. SCALING_POLICY
- D. RESOURCE_CONSTRAINT
- E. ENABLE_QUERY_ACCELERATION

Answer: ([SHOW ANSWER](#))

Snowflake warehouses can use different warehouse types and resource constraints. Snowpark-optimized warehouses use warehouse properties such as `WAREHOUSE_TYPE` and `RESOURCE_CONSTRAINT` to define the compute configuration.

Reference: Snowflake Documentation - Snowpark-optimized warehouses; `CREATE WAREHOUSE`; `ALTER WAREHOUSE`.

NEW QUESTION: 305

Who can access the data published in a Data Exchange?

- A. Only the data provider who published the data
- B. Any Snowflake user, regardless of their role or permissions
- C. Any user who has been provided with a unique Data Exchange URL
- D. Only the users that the data provider has invited to the Data Exchange

Answer: ([SHOW ANSWER](#))

Snowflake's Data Exchange allows data providers to publish data that is accessible only to users they explicitly invite. This mechanism ensures that only approved users or Snowflake accounts have access, adding a layer of security and control. Invited users can query the shared data directly within their Snowflake account without requiring data copies, reducing redundancy and enhancing data accessibility within controlled boundaries. This selective sharing feature is part of Snowflake's secure data-sharing capabilities.

NEW QUESTION: 306

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 307

Which roles can make grant decisions to objects within a managed access schema? (Select TWO)

- A. `ACCOUNTADMIN`
- B. `SECURITYADMIN`
- C. `SYSTEMADMIN`
- D. `ORGADMIN`
- E. `USERADMIN`

Answer: ([SHOW ANSWER](#))

Managed Access Schemas: These are a special type of schema designed for fine-grained access control in Snowflake. Roles with Grant Authority:

ACCOUNTADMIN: The top-level administrative role can grant object privileges on all objects within the account, including managed access schemas.

SECURITYADMIN: Can grant and revoke privileges on objects within the account, including managed access schemas.

Important Note: The ORGADMIN role focuses on organization-level management, not object access control.

NEW QUESTION: 308

Which Snowflake edition enables data sharing only through Snowflake Support?

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

Answer: ([SHOW ANSWER](#))

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 Support.

NEW QUESTION: 309

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: ([SHOW ANSWER](#))

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 process.

NEW QUESTION: 310

Which privilege enables the ability to resume or suspend a task?

- A. MONITOR
- B. OPERATE
- C. CREATE TASK
- D. EXECUTE TASK

Answer: B ([LEAVE A REPLY](#))

The correct answer is B. OPERATE .

The OPERATE privilege on a task allows a role to suspend or resume that task. This is different from privileges used to create, monitor, or execute tasks.

Why B is correct:

To manually suspend or resume a task, the role must have the OPERATE privilege on the task.

Example concept:

GRANT OPERATE ON TASK my_task TO ROLE my_role;

After receiving this privilege, the role can perform actions such as:

ALTER TASK my_task SUSPEND;

ALTER TASK my_task RESUME;

Why the other options are incorrect:

A). MONITOR allows viewing task details and history, but it does not allow suspending or resuming the task.

C). CREATE TASK allows creating tasks in a schema, but it does not automatically allow operating an existing task.

D). EXECUTE TASK is an account-level privilege required for task execution, but it is not the privilege used to suspend or resume a task.

Official Snowflake documentation reference:

Snowflake documentation explains that the OPERATE privilege on a task enables suspending and resuming the task.

Reference: Snowflake Documentation - Task privileges; Snowflake Documentation - GRANT < privileges

> ; SnowPro Core Study Guide - SQL and Snowflake Objects.

NEW QUESTION: 311

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

A. 7 days

B. 14 days

C. 28 days

D. 365 days

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 312

What action will change a customized schedule for a resource monitor back to the default schedule?

A. Changing the monitoring frequency to the default value

B. Changing the monitoring start timestamp to the default value

C. Dropping and re-creating the monitor

D. Cloning the monitor and then editing the cloned monitor schedule

Answer: ([SHOW ANSWER](#))

Resource monitor schedules cannot be reset via update. The only way to restore the default schedule is to drop and recreate the monitor.

Reference:

Snowflake Docs: Resource Monitor Scheduling

NEW QUESTION: 313

Which chart type is supported in Snowsight for Snowflake users to visualize data with dashboards?

A. Area chart

B. Box plot

C. Heat grid

D. Pie chart

Answer: (SHOW ANSWER)

Snowsight, Snowflake 's user interface for exploring, analyzing, and visualizing data, supports a variety of chart types for creating dashboards and visualizations. One of the supported chart types in Snowsight is the Area Chart (A). Area charts are useful for representing quantities over time and can be used to highlight volume change and rate of change, as well as to compare multiple quantities.

While Snowsight supports many types of visualizations to help users analyze their data effectively, including line charts, bar charts, and scatter plots, it 's important to select the specific reference documentation or release notes for the most current list of supported chart types, as Snowflake continues to enhance and update Snowsight 's capabilities.

As of the last update, Box plots (B), Heat grids (C), and Pie charts (D) are types of visualizations that may be supported in various analytics and visualization tools, but for the specific context of Snowsight 's currently confirmed features, Area charts are a verified option for users to visualize their data.

Reference: Snowflake Documentation on Snowsight (<https://docs.snowflake.com/en/user-guide/ui-snowsight.html>)

NEW QUESTION: 314

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: (SHOW ANSWER)

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 temporarybypass 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: 315

When working with a managed access schema, who has the OWNERSHIP privilege of any tables added to the schema?

- A.** The database owner
- B.** The object owner
- C.** The schema owner
- D.** The Snowflake user's role

Answer: (SHOW ANSWER)

In a managed access schema, the schema owner retains the OWNERSHIP privilege of any tables added to the schema. This means that while object owners have certain privileges over the objects they create, only the schema owner can manage privilege grants on these objects¹.

NEW QUESTION: 316

Which constraint is enforced for standard Snowflake tables?

- A. NOT NULL
- B. PRIMARY KEY
- C. FOREIGN KEY
- D. UNIQUE

Answer: ([SHOW ANSWER](#))

The correct answer is A. NOT NULL .

For standard Snowflake tables, Snowflake enforces the NOT NULL constraint. Other constraints such as PRIMARY KEY, FOREIGN KEY, and UNIQUE can be defined, but they are generally informational and are not enforced for standard tables.

Why A is correct:

A NOT NULL constraint prevents NULL values from being inserted into the constrained column.

Example:

```
CREATE TABLE customers (  
customer_id NUMBER NOT NULL,  
customer_name STRING  
);
```

Why the other options are incorrect:

- B). PRIMARY KEY constraints are not enforced for standard Snowflake tables.
- C). FOREIGN KEY constraints are not enforced for standard Snowflake tables.
- D). UNIQUE constraints are not enforced for standard Snowflake tables.

Official Snowflake documentation reference:

Snowflake documentation explains that for standard tables, NOT NULL constraints are enforced, while unique, primary key, and foreign key constraints are not enforced.

Reference: Snowflake Documentation - Constraints; Snowflake Documentation - Table constraints; SnowPro Core Study Guide - SQL and Snowflake Objects.

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 317

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: ([SHOW ANSWER](#))

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: 318

How does Snowflake allow a data provider with an Azure account in central Canada to share data with a data consumer on AWS in Australia?

- A. The data provider in Azure Central Canada can create a direct share to AWS Asia Pacific, if they are both in the same organization.
- B. The data consumer and data provider can form a Data Exchange within the same organization to create a share from Azure Central Canada to AWS Asia Pacific.
- C. The data provider uses the GET DATA workflow in the Snowflake Data Marketplace to create a share between Azure Central Canada and AWS Asia Pacific.
- D. The data provider must replicate the database to a secondary account in AWS Asia Pacific within the same organization then create a share to the data consumer's account.

Answer: ([SHOW ANSWER](#))

Snowflake allows data providers to share data with consumers across different cloud platforms and regions through database replication. The data provider must replicate the database to a secondary account in the target region or cloud platform within the same organization, and then create a share to the data consumer's account. This process ensures that the data is available in the consumer's region and on their cloud platform, facilitating seamless data sharing. References: Sharing data securely across regions and cloud platforms | Snowflake Documentation

NEW QUESTION: 319

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

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

Answer: ([SHOW ANSWER](#))

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

References:

Snowflake Documentation on Reader Accounts

NEW QUESTION: 320

How does Snowflake utilize clustering information to improve query performance?

- A. It prunes unnecessary micro-partitions based on clustering metadata.
- B. It compresses the data within micro-partitions for faster querying.
- C. It automatically allocates additional resources to improve query execution.
- D. It organizes clustering information to speed-up data retrieval from storage

Answer: A ([LEAVE A REPLY](#))

Snowflake utilizes clustering information to enhance query performance by pruning unnecessary micro-partitions.

Clustering Metadata: Snowflake stores clustering information for each micro-partition, which includes data range and distribution.

Pruning Micro-partitions: When a query is executed, Snowflake uses this clustering metadata to identify and eliminate micro-partitions that do not match the query criteria, thereby reducing the amount of data scanned and improving query performance.

Snowflake Documentation on Clustering

Snowflake Documentation on Micro-partition Pruning

NEW QUESTION: 321

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: ([SHOW ANSWER](#)**)**

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-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 322

Which command is used to take away staged files from a Snowflake stage after a successful data ingestion?

- A. DELETE
- B. DROP
- C. REMOVE
- D. TRUNCATE

Answer: ([SHOW ANSWER](#)**)**

The REMOVE command is used in Snowflake to delete files from a stage after they have been successfully ingested into Snowflake tables. This command helps manage storage by allowing users to clean up staged files that are no longer needed, ensuring that the stage does not accumulate unnecessary data over time.

Unlike DELETE, DROP, or TRUNCATE commands, which are used for managing data within Snowflake tables or dropping objects, REMOVE specifically targets the management of files in stages.

References: Snowflake Documentation on Stages and File Management

NEW QUESTION: 323

What will happen to partitions if a table is cloned and the original table type is then changed from permanent to transient?

- A. Both old and new partitions will be changed to transient, but the original table will keep the original table type.
- B. Old partitions will not be changed to transient, but any new partitions added to the clone will follow the transient lifecycle.
- C. Both old and new partitions will be changed to transient, and the original table type will change to transient.
- D. Old partitions will be changed to transient, and any new partitions added to the clone will follow the transient lifecycle.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 324

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

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

Answer: ([SHOW ANSWER](#))

Micro-partitions in Snowflake are immutable objects, which means once they are written, they cannot be modified. This immutability supports the use of Time Travel, allowing users to access historical data within a defined period.

Additionally, micro-partitions can significantly reduce the amount of I/O from object storage to virtual warehouses. This is because Snowflake's query optimizer can skipover micro-partitions that do not contain relevant data for a query, thus reducing the amount of data that needs to be scanned and transferred.

References: [COF-C03] SnowPro Core Certification Exam Study Guide

<https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions.html>

NEW QUESTION: 325

What action should be taken if a large number of concurrent queries are queued in a virtual warehouse?

- A. Scale-up by resizing the warehouse.
- B. Disable auto-suspend on the warehouse.
- C. Scale-out with a multi-cluster warehouse.
- D. Enable auto-resume on the warehouse.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 326

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

- A. Load files that are approximately 25 MB or smaller.

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

Answer: ([SHOW ANSWER](#))

When loading data into Snowflake, it is recommended to:

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

D). Avoid using embedded characters such as commas for numeric data types: Embedded characters can cause issues during data loading as they may be interpreted incorrectly. It's best to clean the data of such characters to ensure accurate and efficient data loading.

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

References:

Snowflake Documentation on DataLoading Considerations

[COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 327

Cloning is being used to create a development environment. Will storage be impacted when cloned tables are modified in a development environment?

- A. Cloned tables will be duplicated in storage.
- B. Storage consumption remains unaffected.
- C. The modified micro-partitions consume additional storage.
- D. Storage usage increases with the number of clones created.

Answer: ([SHOW ANSWER](#))

Zero-copy cloning initially shares micro-partitions between the source object and the clone. However, when data is modified in either the original or cloned object, Snowflake creates new micro-partitions for the changed data. Those modified micro-partitions consume additional storage.

Reference: Snowflake Documentation - Zero-copy cloning; Micro-partitions; Storage impact of cloned objects.

NEW QUESTION: 328

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

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

Answer: ([SHOW ANSWER](#))

TheSYSTEM\$GET_PRIVATELINKfunction in Snowflake provides proxy information necessary for configuring PrivateLink connections, which can protect Snowsight as well as other Snowflake services.

PrivateLink enhances security by allowing Snowflake to be accessed via a private connection within a cloud provider's network, reducing exposure to the public internet.

References:

Snowflake Documentation: PrivateLink Setup

NEW QUESTION: 329

Which statement MOST accurately describes clustering in Snowflake?

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

Answer: ([SHOW ANSWER](#))

Clustering in Snowflake refers to the organization of data within micro-partitions, which are contiguous units of storage within Snowflake tables. Clustering keys can be defined to co-locate similar rows in the same micro-partitions, improving scan efficiency and query performance¹².

References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 330

Which minimum Snowflake edition allows for a dedicated metadata store?

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

Answer: ([SHOW ANSWER](#))

The Enterprise edition of Snowflake allows for a dedicated metadata store, providing additional features designed for large-scale enterprises Reference: <https://docs.snowflake.com/en/user-guide/intro-editions.html>

NEW QUESTION: 331

Which action will scale-out a Snowflake virtual warehouse?

- A. Adding 2 additional clusters to a multi-cluster warehouse
- B. Removing 1 cluster from a multi-cluster warehouse
- C. Resizing a warehouse from size Small to X-Large
- D. Resizing a warehouse from size Large to Medium

Answer: ([SHOW ANSWER](#))

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here: <https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 332

If a multi-cluster virtual warehouse is running in Auto-scale mode, how are workloads with unpredictable usage patterns handled?

- A. Snowflake automatically resizes the warehouse based on usage requirements.
- B. Snowflake automatically starts and stops additional clusters as needed.
- C. Snowflake notifies users when to adjust the number of clusters to improve workload performance.
- D. Snowflake automatically starts additional clusters when the primary cluster reaches 50% capacity.

Answer: (SHOW ANSWER)

In Auto-scale mode, a multi-cluster virtual warehouse can automatically scale out by starting additional clusters as needed to handle concurrent queries . This is ideal for unpredictable or bursty workloads.

- * Clusters spin up based on queue length and system load .
- * They automatically shut down when not in use .

Reference:

Snowflake Docs: Multi-Cluster Warehouses

NEW QUESTION: 333

How does the Access_History view enhance overall data governance pertaining to read and write operations? (Select TWO).

- A. Shows how the accessed data was moved from the source to the target objects
- B. Provides a unified picture of what data was accessed and when it was accessed
- C. Protects sensitive data from unauthorized access while allowing authorized users to access it at query runtime
- D. Identifies columns with personal information and tags them so masking policies can be applied to protect sensitive data
- E. Determines whether a given row in a table can be accessed by the user by filtering the data based on a given policy

Answer: (SHOW ANSWER)

The ACCESS_HISTORY view in Snowflake is a powerful tool for enhancing data governance, especially concerning monitoring and auditing data access patterns for both read and write operations. The key ways in which ACCESS_HISTORY enhances overall data governance are:

- B). Provides a unified picture of what data was accessed and when it was accessed: This view logs details about query executions, including the objects (tables, views) accessed and the timestamps of these accesses. It ' s instrumental in auditing and compliance scenarios, where understanding the access patterns to sensitive data is critical.
- E). Determines whether a given row in a table can be accessed by the user by filtering the data based on a given policy: While this option is a bit of a misinterpretation of what ACCESS_HISTORY directly offers, it indirectly supports data governance by providing the information necessary to analyze access patterns. This analysis can then inform policy decisions, such as implementing Row-Level Security (RLS) to restrict access to specific rows based on user roles or attributes.

ACCESS_HISTORY does not automatically apply data masking or tag columns with personal information.

However, the insights derived from analyzing ACCESS_HISTORY can be used to identify sensitive data and inform the application of masking policies or other security measures.

References:

Snowflake Documentation on ACCESS_HISTORY: Access History

NEW QUESTION: 334

A permanent table and temporary table have the same name, TBL1, in a schema.

What will happen if a user executes `select * from TBL1 ;`?

- A. The temporary table will take precedence over the permanent table.
- B. The permanent table will take precedence over the temporary table.
- C. An error will say there cannot be two tables with the same name in a schema.
- D. The table that was created most recently will take precedence over the older table.

Answer: ([SHOW ANSWER](#))

In Snowflake, if a temporary table and a permanent table have the same name within the same schema, the temporary table takes precedence over the permanent table within the session where the temporary table was created⁴.

NEW QUESTION: 335

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

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

Answer: ([SHOW ANSWER](#))

Snowpark is designed to bring the data programmability to Snowflake, enabling developers to write code in familiar languages like Scala, Java, and Python. It allows for the execution of these codes directly within Snowflake's virtual warehouses, eliminating the need for a separate cluster. Additionally, Snowpark's compatibility with Spark allows users to leverage their existing Spark code with minimal changes¹.

NEW QUESTION: 336

What is the MINIMUM Snowflake edition required to add masking policies to selectively mask plain-text data in a table or in view columns at query time?

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

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 337

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: (SHOW ANSWER)

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#:~:](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.)

text=The%20parameter%20STATEMENT_QUEUED_TIMEOUT_IN_SECONDS%20sets%20the, indefinitely%20in %20the%20waiting%20queue.

NEW QUESTION: 338

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: (SHOW ANSWER)

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: 339

Which governance feature is supported by all Snowflake editions?

- A. Object tags
- B. Masking policies
- C. Row access policies
- D. OBJECT_DEPENDENCIES View

Answer: (SHOW ANSWER)

Snowflake 's governance features vary across different editions, but the OBJECT_DEPENDENCIES view is supported by all Snowflake editions. This feature is part of Snowflake 's Information Schema and is designed to help users understand the dependencies between various objects in their Snowflake environment.

The OBJECT_DEPENDENCIES view provides a way to query and analyze the relationships and dependencies among different database objects, such as tables, views, and stored procedures. This is crucial for governance, as it allows administrators and data engineers to assess the impact of changes, understand object relationships, and ensure proper management of data assets.

Object tags, masking policies, and row access policies are more advanced features that offer fine-grained data governance capabilities such as tagging objects for classification, dynamically masking sensitive data based on user

roles, and controlling row-level access to data. These features may have varying levels of support across different Snowflake editions, with some features being exclusive to higher-tier editions.

Reference: Snowflake Documentation on Information Schema (https://docs.snowflake.com/en/sql-reference/info-schema/object_dependencies.html)

NEW QUESTION: 340

Which Snowflake partner specializes in data catalog solutions?

- A. Alation
- B. DataRobot
- C. dbt
- D. Tableau

Answer: ([SHOW ANSWER](#))

Alation is known for specializing in data catalog solutions and is a partner of Snowflake. Data catalog solutions are essential for organizations to effectively manage their metadata and make it easily accessible and understandable for users, which aligns with the capabilities provided by Alation.

References:

[COF-C03] SnowPro Core Certification Exam Study Guide
Snowflake's official documentation and partner listings

NEW QUESTION: 341

Which virtual warehouse auto-suspend configurations will result in a warehouse that runs continually 24 hours a day, 7 days a week? (Select TWO).

- A. -1
- B. 0
- C. NULL
- D. 86400
- E. 172800

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 342

Which Snowflake database object can be shared with other accounts?

- A. Tasks
- B. Pipes
- C. Secure User-Defined Functions (UDFs)
- D. Stored Procedures

Answer: ([SHOW ANSWER](#))

In Snowflake, Secure User-Defined Functions (UDFs) can be shared with other accounts using Snowflake's data sharing feature. This allows different Snowflake accounts to securely execute the UDFs without having direct access to the underlying data the functions operate on, ensuring privacy and security. The sharing is facilitated through shares created in Snowflake, which can contain Secure UDFs along with other database objects like tables and views.

References: Snowflake Documentation on Data Sharing and Secure UDFs

NEW QUESTION: 343

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: ([SHOW ANSWER](#))

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-C03] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

NEW QUESTION: 344

Which URL type should be used to provide a user with a window of time to access files that are in cloud storage?

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

Answer: ([SHOW ANSWER](#))

The correct answer is C. Pre-signed URL .

Corrected wording:

The option listed as "re-signed URL" appears to be a typing error. The correct term is "pre-signed URL." A pre-signed URL provides temporary access to a file in cloud storage. It is commonly used when a user or application needs time-limited access to a staged file without granting broader permissions to the underlying storage location.

Why C is correct:

A pre-signed URL includes authorization information and is valid only for a defined period of time. This makes it appropriate when users need a limited access window.

Example concept:

```
SELECT GET_PRESIGNED_URL(@my_stage, ' path/file.csv ', 3600);
```

This returns a URL that allows access to the staged file for a limited duration, such as 3600 seconds.

Why the other options are incorrect:

- A). A scoped URL is encoded and scoped to a query result context. It is useful for controlled access through Snowflake, but the option that directly provides a time-limited access window to cloud storage is a pre-signed URL.
- B). A file URL identifies a file location, but it does not itself provide temporary authorized access.
- D). Redirect URL is not the Snowflake file access URL type used for this requirement.

Official Snowflake documentation reference:

Snowflake documentation describes GET_PRESIGNED_URL as a function that returns a pre-signed URL to a staged file, allowing temporary access for a specified expiration time.

Reference: Snowflake Documentation - GET_PRESIGNED_URL; Snowflake Documentation - File functions; SnowPro Core Study Guide - Data Loading and Unloading.

NEW QUESTION: 345

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

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

Answer: ([SHOW ANSWER](#)**)**

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

NEW QUESTION: 346

Which queries will require a running virtual warehouse? (Select TWO).

- A. UPDATE < table > SET;
- B. SELECT COUNT(*) FROM < table > ;
- C. DELETE FROM < table > ;
- D. SELECT CURRENT_DATABASE();
- E. DROP TABLE < table > ;

Answer: ([SHOW ANSWER](#)**)**

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 347

What information in a secure User-Defined Function, or UDF, is only available to authorized users? Select TWO.

- A. The UDF definition code
- B. Return type
- C. Parameter types
- D. List of data imports
- E. Handler language

Answer: ([SHOW ANSWER](#)**)**

The correct answers are A. The UDF definition code and D. List of data imports .

A secure UDF is designed to protect sensitive implementation details. Secure UDFs help prevent unauthorized users from viewing the function body and other details that could expose business logic or sensitive references.

Why A is correct:

The definition code of a secure UDF is hidden from unauthorized users. This protects the logic used inside the function.

Why D is correct:

The list of imported files or data imports associated with the secure UDF is also protected from unauthorized users because it may reveal sensitive implementation details.

Why the other options are incorrect:

B). The return type is generally part of the function signature and must be visible so users can understand how to call and use the function.

C). Parameter types are also part of the function signature and are visible to users who need to call the UDF.

E). The handler language may be visible as part of function metadata and is not the best answer for protected secure UDF information.

Official Snowflake documentation reference:

Snowflake documentation explains that secure UDFs hide internal details, including the function definition and sensitive implementation information, from unauthorized users.

Reference: Snowflake Documentation - Secure UDFs; Snowflake Documentation - User-defined functions; SnowPro Core Study Guide - Security and Access Control.

NEW QUESTION: 348

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

A. All child objects will retain data for the new retention period.

B. All child objects that do not have an explicit retention period will automatically inherit the new retention period.

C. All child objects with an explicit retention period will be overridden with the new retention period.

D. All explicit child object retention periods will remain unchanged.

Answer: (SHOW ANSWER)

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

NEW QUESTION: 349

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

A. STRUCT

B. OBJECT

C. BINARY

D. ARRAY

E. CLOB

Answer: (SHOW ANSWER)

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

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

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

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

References:

NEW QUESTION: 350

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: ([SHOW ANSWER](#))

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

NEW QUESTION: 351

Which tasks are performed in the Snowflake Cloud Services layer? (Choose two.)

- A. Management of metadata
- B. Computing the data
- C. Maintaining Availability Zones
- D. Infrastructure security
- E. Parsing and optimizing queries

Answer: ([SHOW ANSWER](#))

The Snowflake Cloud Services layer performs a variety of tasks, including the management of metadata and the parsing and optimization of queries. This layer is responsible for coordinating activities across Snowflake, including user session management, security, and query compilation³.

NEW QUESTION: 352

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 353

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: ([SHOW ANSWER](#))

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-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 354

While unloading data into a stage, how can the user ensure that the output will be a single file?

- A. Use the copy option files=single.
- B. Use the COPY Option SINGLE=TRUE .
- C. Use the get option SINGLE-TRUE.
- D. Use the GET option FILES-SINGLE.

Answer: B ([LEAVE A REPLY](#))

To ensure that the output will be a single file when unloading data into a stage, you should use the COPY option SINGLE=TRUE. This option specifies that the result of the COPY INTO command should be written to a single file, rather than multiple files.

Snowflake Documentation: COPY INTO < location >

NEW QUESTION: 355

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 356

What does the :: operator do when used with a VARIANT value?

- A. Converts a VARIANT value into an object.
- B. Extracts the raw value from the VARIANT column.
- C. Casts the value from a VARIANT to a specified data type.
- D. Generates a random sample from the VARIANT column.

Answer: ([SHOW ANSWER](#))

The correct answer is C. Casts the value from a VARIANT to a specified data type .

In Snowflake, the :: operator is the cast operator. It converts a value from one data type to another. When working with semi-structured data, it is commonly used to cast values extracted from a VARIANT column into standard SQL data types such as STRING, NUMBER, BOOLEAN, or TIMESTAMP.

Why C is correct:

A VARIANT value often needs to be explicitly cast before it is used in comparisons, calculations, or formatted output.

Example:

```
SELECT src:customer_id::NUMBER,  
src:name::STRING  
FROM my_table;
```

In this example:

```
src:customer_id::NUMBER
```

extracts customer_id from a VARIANT column and casts it to NUMBER.

Why the other options are incorrect:

A). :: does not specifically convert every VARIANT into an object. It casts to whichever data type is specified.

B). Extracting a value from a VARIANT uses path notation or functions such as GET and GET_PATH; :: performs casting.

D). :: has nothing to do with random sampling.

Official Snowflake documentation reference:

Snowflake documentation describes :: as the cast operator and shows its use for converting semi-structured VARIANT values to explicit SQL data types.

Reference: Snowflake Documentation - Data type conversion; Snowflake Documentation - Cast operator; Snowflake Documentation - Querying semi-structured data; SnowPro Core Study Guide - Working with Semi-Structured Data.

NEW QUESTION: 357

Which Snowflake edition supports a dedicated metadata store and separate compute clusters that are inaccessible to other Snowflake customers?

A. Standard

B. Enterprise

C. Business Critical

D. Virtual Private Snowflake, or VPS

Answer: D (LEAVE A REPLY)

The correct answer is D. Virtual Private Snowflake, or VPS .

Virtual Private Snowflake is the Snowflake offering designed for organizations with strict isolation requirements. It provides dedicated Snowflake resources, including an isolated metadata store and dedicated compute resources that are not shared with other Snowflake customers.

Why D is correct:

VPS provides the highest level of isolation among the listed Snowflake editions. It is intended for customers with enhanced security, compliance, or regulatory requirements.

Why the other options are incorrect:

- A). Standard Edition is the base Snowflake edition and does not provide dedicated isolated metadata and compute clusters.
- B). Enterprise Edition adds features such as multi-cluster warehouses and extended Time Travel, but it does not provide VPS-level isolation.
- C). Business Critical Edition provides enhanced security and compliance features, but not the dedicated isolated Snowflake environment described in the question.

Official Snowflake documentation reference:

Snowflake documentation describes Virtual Private Snowflake as offering dedicated resources and isolation, including dedicated metadata services and compute resources inaccessible to other customers.

Reference: Snowflake Documentation - Snowflake editions; Virtual Private Snowflake; SnowPro Core Study Guide - Snowflake Architecture.

NEW QUESTION: 358

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 359

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 360

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

- A. 1 minute
- B. 2 minutes

- C. 6 minutes
- D. 8 minutes

Answer: B ([LEAVE A REPLY](#))

In a multi-cluster warehouse with an economy scaling policy, queries will wait in the queue for 2 minutes before another cluster is started. This is to minimize costs by allowing queries to queue up for a short period before adding additional compute resources. References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 361

How can the outer array structure of a semi-structured file be removed?

- A. Use the parameter `strip_outer_array = true` in a `COPY INTO <table>` command.
- B. Set the file format to eliminate any outer array structure before initiating the `COPY INTO <table>` command.
- C. Filter the outer array structure using a `PUT` command with the `include_outer_array = false` parameter.
- D. Use the `FLATTEN` command with the `outer_array = false` parameter.

Answer: A ([LEAVE A REPLY](#))

In Snowflake, the parameter `strip_outer_array = true` can be set in the `COPY INTO <table>` command to remove the outer array structure from a semi-structured file. This parameter is useful for JSON files or similar data formats where an array might wrap the data, allowing Snowflake to directly load the inner elements as table rows.

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:

<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount

Code: **freecram**)

NEW QUESTION: 362

Which resource monitor setting will cancel all active queries in a virtual warehouse when the threshold is met?

- A. SUSPEND
- B. NOTIF
- C. NOTIFY_USERS
- D. SUSPEND_IMMEDIATE

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 363

Which data sharing mechanism can be used to share data privately or publicly within the Snowflake Marketplace?

- A. Listing
- B. Reader account
- C. Data Exchange
- D. Direct share

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 364

When would a user use the GET command?

- A. A staged file needs to be unloaded to a local machine.
- B. A file on a local machine needs to be loaded onto a stage.
- C. A query result needs to be opened in a third-party tool.
- D. A file needs to be permanently deleted from Snowflake.

Answer: (SHOW ANSWER)

The correct answer is A. A staged file needs to be unloaded to a local machine .

The Snowflake GET command is used to download files from an internal Snowflake stage to a local file system. It is commonly used after data has been unloaded from a table into an internal stage and the user wants to retrieve those staged files locally.

Why A is correct:

The GET command downloads files from a Snowflake internal stage to a local directory.

Example:

```
GET @my_stage file:///tmp/data/;
```

This command downloads files from the internal stage @my_stage to the local /tmp/data/ directory.

Why the other options are incorrect:

- B). Loading a local file onto a stage is done with the PUT command, not GET.
- C). Opening query results in a third-party tool is not the purpose of the GET command.
- D). Permanently deleting files from a Snowflake stage is done with the REMOVE command, not GET.

Official Snowflake documentation reference:

Snowflake documentation describes the GET command as the command used to download data files from a Snowflake stage to a local directory or folder.

Reference: Snowflake Documentation - GET command; Snowflake Documentation - PUT command; Snowflake Documentation - REMOVE command; SnowPro Core Study Guide - Data Loading and Unloading.

NEW QUESTION: 365

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: (SHOW ANSWER)

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: 366

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 367

Which copy INTO command outputs the data into one file?

- A. SINGLE=TRUE
- B. MAX_FILE_NUMBER=1
- C. FILE_NUMBER=1
- D. MULTIPLE=FAISE

Answer: B ([LEAVE A REPLY](#))

The COPY INTO command in Snowflake can be configured to output data into a single file by setting the MAX_FILE_NUMBER option to 1. This option limits the number of files generated by the command, ensuring that only one file is created regardless of the amount of data being exported.

References:

[COF-C03] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Data Unloading

NEW QUESTION: 368

What is a recommended best practice when loading data into Snowflake?

- A. Combine smaller files into larger tiles before loading.
- B. When loading a single large file split it into multiple files before loading.
- C. Set the DATA_RETENTION_TIME_IN_DAYS on the table to 0 before loading.
- D. Make sure the files in the stage are either all compressed or all uncompressed.

Answer: ([SHOW ANSWER](#))

To optimize parallelism and maximize performance, large files should be split into multiple smaller files (ideally 100-250 MB compressed) before loading. This allows Snowflake to load files in parallel using multiple threads or clusters.

Reference:

Snowflake Docs: Best Practices for Data Loading

NEW QUESTION: 369

Which data formats are supported by Snowflake when unloading semi-structured data? (Select TWO).

- A. Binary file in Avro
- B. Binary file in Parquet
- C. Comma-separated JSON
- D. Newline Delimited JSON
- E. Plain text file containing XML elements

Answer: ([SHOW ANSWER](#))

Snowflake supports a variety of file formats for unloading semi-structured data, among which Parquet and Newline Delimited JSON (NDJSON) are two widely used formats.

B). Binary file in Parquet: Parquet is a columnar storage file format optimized for large-scale data processing and analysis. It is especially suited for complex nested data structures.

D). Newline Delimited JSON (NDJSON): This format represents JSON records separated by newline characters, facilitating the storage and processing of multiple, separate JSON objects in a single file.

These formats are chosen for their efficiency and compatibility with data analytics tools and ecosystems, enabling seamless integration and processing of exported data.

References:

Snowflake Documentation: Data Unloading

NEW QUESTION: 370

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: ([SHOW ANSWER](#))

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: 371

Taking which step will minimize the number of queries waiting to run on a virtual warehouse?

- A. Increase the WAREHOUSE_SIZE.
- B. Increase the value of the AUTO_SUSPEND parameter.
- C. Set the WAIT_FOR_COMPLETION parameter to FALSE.
- D. Increase the value of the MAX_CLUSTER_COUNT.

Answer: ([SHOW ANSWER](#))

The correct answer is D. Increase the value of the MAX_CLUSTER_COUNT .

Queries wait in a queue when a warehouse does not have enough available compute capacity to handle the current concurrency. For a multi-cluster warehouse, increasing MAX_CLUSTER_COUNT allows Snowflake to start more clusters when concurrency increases.

Why D is correct:

Increasing MAX_CLUSTER_COUNT allows a multi-cluster warehouse in Auto-scale mode to add more clusters, up to the configured maximum. This helps reduce query queuing for concurrent workloads.

Why the other options are incorrect:

- A). Increasing WAREHOUSE_SIZE gives each cluster more compute power, which can help individual queries run faster, but it is not the best direct solution for reducing queued concurrent queries.
- B). Increasing AUTO_SUSPEND keeps the warehouse running longer but does not increase concurrency capacity.
- C). WAIT_FOR_COMPLETION = FALSE is used with warehouse resize behavior and does not reduce query queuing.

Official Snowflake documentation reference:

Snowflake documentation explains that multi-cluster warehouses improve concurrency by starting additional clusters as query demand increases. The MAX_CLUSTER_COUNT property defines the maximum number of clusters that can run.

Reference: Snowflake Documentation - Multi-cluster warehouses; Snowflake Documentation - Warehouse properties; SnowPro Core Study Guide - Snowflake Account and Resource Management.

NEW QUESTION: 372

What is the default authentication method while using the JDBC driver connection in Snowflake?

- A. externalbrowser
- B. snowflake
- C. username_password_mfa
- D. snowflake_jwt

Answer: (SHOW ANSWER)

The default authentication mechanism for Snowflake clients (including JDBC) is the snowflake method, which uses username and password authentication.

Reference:

Snowflake Docs: JDBC Driver - Authentication

NEW QUESTION: 373

What feature specifies which roles are allowed to see a particular record in a table?

- A. Dynamic Data Masking
- B. Secure Data Sharing
- C. Row access policy
- D. Object tagging

Answer: (SHOW ANSWER)

The correct answer is C. Row access policy .

A row access policy is used to implement row-level security in Snowflake. It controls whether a row is visible in query results based on conditions such as the current role, current user, or values in a mapping table.

Why C is correct:

The question asks which feature controls visibility of particular records, meaning rows. Row access policies are designed exactly for that purpose.

Example concept:

```
CREATE ROW ACCESS POLICY region_policy
AS (region STRING) RETURNS BOOLEAN - >
region IN (
SELECT allowed_region
FROM role_region_mapping
WHERE role_name = CURRENT_ROLE()
);
```

This policy can allow users with different roles to see only the rows assigned to their region.

Why the other options are incorrect:

- A). Dynamic Data Masking controls how column values are displayed. It masks data in columns but does not filter entire rows.
- B). Secure Data Sharing allows data to be shared between accounts, but it is not the row-level filtering feature itself.
- D). Object tagging classifies or labels objects for governance, but it does not directly control row visibility.

Official Snowflake documentation reference:

Snowflake documentation describes row access policies as schema-level objects that determine whether a row is visible in query results.

Reference: Snowflake Documentation - Row access policies; Snowflake Documentation - Row-level security; SnowPro Core Study Guide - Data Protection and Governance.

NEW QUESTION: 374

A Snowflake user has been granted the create data EXCHANGE listing privilege with their role.

Which tasks can this user now perform on the Data Exchange? (Select TWO).

- A. Rename listings.
- B. Delete provider profiles.
- C. Modify listings properties.
- D. Modify incoming listing access requests.
- E. Submit listings for approval/publishing.

Answer: ([SHOW ANSWER](#))

With the create data EXCHANGE listing privilege, a Snowflake user can modify the properties of listings and submit them for approval or publishing on the Data Exchange. This allows them to manage and share data sets with consumers effectively. References: Based on general data exchange practices in cloud services as of 2021.

NEW QUESTION: 375

What takes the highest precedence in Snowflake file format options, when specified in multiple locations during data loading?

- A. The stage definition
- B. The table definition

- C. The use of a copy into < table > statement
- D. The use of a copy INTO < location > statement

Answer: ([SHOW ANSWER](#))

When loading data into Snowflake, the file format options specified in the COPY INTO < table > statement take the highest precedence over other locations such as the stage or table definitions. This ensures that any specific settings for a particular load operation are applied correctly.

File Format Hierarchy:

Stage Definition: Specifies default file format options for files staged in the location.

Table Definition: Can specify default file format options associated with the table.

COPY INTO Statement: Overrides both the stage and table definitions with the file format options specified directly in the statement.

Example Usage:

```
COPY INTO my_table
```

```
FROM @my_stage
```

```
FILE_FORMAT = (FORMAT_NAME = ' my_format ' FIELD_OPTIONALLY_ENCLOSED_BY = ' " ' ');
```

 References:

Snowflake Documentation: Copy into Table

Snowflake Documentation: File Format Options

NEW QUESTION: 376

Which Snowflake view is used to support compliance auditing?

- A. ACCESS_HISTORY
- B. COPY_HISTORY
- C. QUERY_HISTORY
- D. ROW ACCESS POLICIES

Answer: ([SHOW ANSWER](#))

The ACCESS_HISTORY view in Snowflake is utilized to support compliance auditing. It provides detailed information on data access within Snowflake, including reads and writes by user queries. This view is essential for regulatory compliance auditing as it offers insights into the usage of tables and columns, and maintains a direct link between the user, the query, and the accessed data¹.

References: [COF-C03] SnowPro Core Certification Exam Study Guide

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 377

A marketing co-worker has requested the ability to change a warehouse size on their medium virtual warehouse called mktg__WH.

Which of the following statements will accommodate this request?

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

Answer: B ([LEAVE A REPLY](#))

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

References:

[COF-C03] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Access Control Privileges¹

NEW QUESTION: 378

A table, STUDENTS, is created using this command:

An example value for the marks is:

Which query will return all of the students ' names and history marks?

- A. SELECT NAME, MARKS:history FROM STUDENTS;
- B. SELECT NAME, MARKS: HISTORY FROM STUDENTS;
- C. SELECT NAME, MARKS [0] FROM STUDENTS;
- D. SELECT NAME, GET (MARKS, ' history ') FROM STUDENTS;

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 379

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

- A. True
- B. False

Answer: ([SHOW ANSWER](#)**)**

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: 380

Which parameter prevents streams on tables from becoming stale?

- A. MAXDATAEXTENSIONTIMEINDAYS
- B. MTN_DATA_RETENTION_TTIME_TN_DAYS

C. LOCK_TIMEOUT

D. STALE_AFTER

Answer: ([SHOW ANSWER](#))

The parameter that prevents streams on tables from becoming stale is

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

NEW QUESTION: 381

Which Snowflake function and command combination should be used to convert rows in a relational table to a single VARIANT column, and unload the rows into a file in JSON format? (Select TWO).

A. PUT

B. GET

C. COPY

D. EXPORT

E. OBJECT_CONSTRUCT

Answer: ([SHOW ANSWER](#))

To convert rows in a relational table to a single VARIANT column and unload the rows into a file in JSON format, you can use the COPY command in combination with the OBJECT_CONSTRUCT function. The OBJECT_CONSTRUCT function converts the row into a JSON object stored in a VARIANT column, and the COPY command can then be used to unload this data into a JSON file.

Snowflake Documentation: OBJECT_CONSTRUCT

Snowflake Documentation: COPY INTO < location >

Top of Form

Bottom of Form

NEW QUESTION: 382

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

A. UPDATE TABLE FROM STREAM

B. SELECT FROM STREAM

C. INSERT INTO TABLE SELECT FROM STREAM

D. ALTER TABLE AS SELECT FROM STREAM

E. BEGIN COMMIT

Answer: ([SHOW ANSWER](#))

The SQL commands that consume a stream and advance the stream offset are those that result in changes to the data, such as UPDATE and INSERT operations. Specifically, 'UPDATE TABLE FROM STREAM' and 'INSERT INTO TABLE SELECT FROM STREAM' will consume the stream and move the offset forward, reflecting the changes made to the data.

References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 383

What are characteristics of reader accounts in Snowflake? (Select TWO).

- A. Reader account users cannot add new data to the account.
- B. Reader account users can share data to other reader accounts.
- C. A single reader account can consume data from multiple provider accounts.
- D. Data consumers are responsible for reader account setup and data usage costs.
- E. Reader accounts enable data consumers to access and query data shared by the provider.

Answer: ([SHOW ANSWER](#))

Characteristics of reader accounts in Snowflake include:

A . Reader account users cannot add new data to the account: Reader accounts are intended for data consumption only. Users of these accounts can query and analyze the data shared with them but cannot upload or add new data to the account.

E . Reader accounts enable data consumers to access and query data shared by the provider: One of the primary purposes of reader accounts is to allow data consumers to access and perform queries on the data shared by another Snowflake account, facilitating secure and controlled data sharing.

References:

Snowflake Documentation: Reader Accounts

NEW QUESTION: 384

What does a table with a clustering depth of 1 mean in Snowflake?

- A. The table has only 1 micro-partition.
- B. The table has 1 overlapping micro-partition.
- C. The table has no overlapping micro-partitions.
- D. The table has no micro-partitions.

Answer: ([SHOW ANSWER](#))

In Snowflake, a table ' s clustering depth indicates the degree of micro-partition overlap based on the clustering keys defined for the table. A clustering depth of 1 implies that the table has no overlapping micro- partitions. This is an optimal scenario, indicating that the table ' s data is well-clustered according to the specified clustering keys. Well-clustered data can lead to more efficient query performance, as it reduces the amount of data scanned during query execution and improves the effectiveness of data pruning.

References:

Snowflake Documentation on Clustering: Understanding Clustering Depth

NEW QUESTION: 385

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: ([SHOW ANSWER](#))

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: 386

Which copy option is used to include column headers when unloading data to Parquet files?

- A. PARSE_HEADER = TRUE
- B. SKIP_HEADER = < integer >
- C. HEADER = TRUE
- D. HEADER = (' < header_1 > ' = ' < value_1 > ')

Answer: ([SHOW ANSWER](#))

The correct answer is C. HEADER = TRUE .

When unloading data with the COPY INTO < location > command, the HEADER option specifies whether column names should be included in the output file.

Why C is correct:

HEADER = TRUE includes column headings in the unloaded file output.

Example:

```
COPY INTO @my_stage/output/  
FROM my_table  
FILE_FORMAT = (TYPE = PARQUET)  
HEADER = TRUE;
```

Important note for exam review:

For CSV unloads, HEADER = TRUE includes column names in the first row. For Parquet unloads, Snowflake can use column names in the Parquet output metadata/schema when the header option is enabled.

Why the other options are incorrect:

- A). PARSE_HEADER is a file format option used when loading CSV data with inferred column names. It is not the unload option for including headers.
- B). SKIP_HEADER is used when loading files to skip header rows. It does not include headers during unloading.
- D). This syntax is not the correct Snowflake unload option for column headers.

Official Snowflake documentation reference:

Snowflake documentation for COPY INTO < location > describes the HEADER option as specifying whether to include table column headings in the output files.

Reference: Snowflake Documentation - COPY INTO < location > ; Snowflake Documentation - Unloading data;

SnowPro Core Study Guide - Data Loading and Unloading.

NEW QUESTION: 387

What would cause different results to be returned when running the same query twice?

- A. SAMPLE is used and the seed is set
- B. sample is used and the seed is not set.
- C. Fraction-based sampling is used.
- D. Fixed-size sampling is used.

Answer: ([SHOW ANSWER](#))

When using the SAMPLE clause in a query, if the seed is not set, Snowflake will use a different random seed for each execution of the query. This results in different rows being sampled each time, leading to different results. Setting a seed ensures that the same rows are sampled each time the query is run.

Snowflake Documentation: Sampling

NEW QUESTION: 388

Which items are considered schema objects in Snowflake? (Select TWO).

- A. Pipe
- B. File format
- C. Resource monitor
- D. Storage integration
- E. Virtual warehouse

Answer: ([SHOW ANSWER](#))

In Snowflake, schema objects include Pipes and File formats. Pipes are used for continuous data loading, and File formats specify the format of data files used in loading and unloading operations within Snowflake

NEW QUESTION: 389

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

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

Answer: ([SHOW ANSWER](#))

The query acceleration service in Snowflake is designed to benefit queries that involve large scans and selective filters. This service can offload portions of the query processing work to shared compute resources, which can handle these types of workloads more efficiently by performing more work in parallel and reducing the wall-clock time spent in scanning and filtering. References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 390

What is the primary purpose for configuring multiple Identity Providers, or IdPs, in Snowflake?

- A. To enforce complex password rules
- B. To improve the performance of a query
- C. To manage encryption keys for secure data storage
- D. To allow different authentication methods for multiple user groups

Answer: ([SHOW ANSWER](#))

The correct answer is D. To allow different authentication methods for multiple user groups .

Multiple identity providers allow an organization to support different authentication configurations for different users, departments, partners, or applications. This is useful when separate user groups authenticate through different SSO providers or require different authentication flows.

Why D is correct:

Configuring multiple IdPs allows Snowflake to support multiple federated authentication integrations and assign appropriate authentication behavior to different users or groups.

Why the other options are incorrect:

- A). Complex password rules are controlled by password policies, not by configuring multiple IdPs.
- B). Identity providers affect authentication, not query performance.
- C). Encryption key management is handled by Snowflake encryption features and key management options, not IdP configuration.

Official Snowflake documentation reference:

Snowflake documentation describes federated authentication and multiple identity provider support as a way to configure different SSO/authentication integrations for different users.

Reference: Snowflake Documentation - Federated authentication; Multiple identity providers; Authentication policies; SnowPro Core Study Guide - Security and Access Control.

NEW QUESTION: 391

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: (SHOW ANSWER)

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 objects.

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 392

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

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

Answer: D,E (LEAVE A REPLY)

Semi-structured
JSON, Parquet

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

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

NEW QUESTION: 393

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: ([SHOW ANSWER](#))

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: 394

Which function should be used to access values from an object?

- A. OBJECT_KEYS
- B. OBJECT_PICK
- C. GET_PATH
- D. XMLGET

Answer: ([SHOW ANSWER](#))

The correct answer is C. GET_PATH .

GET_PATH extracts a value from semi-structured data using a path expression. It can be used to access values inside an object, including nested values.

Why C is correct:

For an object stored in a VARIANT or OBJECT column, GET_PATH can retrieve a value by specifying the key path.

Example:

```
SELECT GET_PATH(src, ' customer.name ' )  
FROM my_table;
```

This accesses the name value inside the customer object.

Why the other options are incorrect:

- A). OBJECT_KEYS returns the keys from an object, not the values for a specific key path.
- B). OBJECT_PICK returns a new object containing selected key-value pairs, but it is not the general function for accessing a value by path.
- D). XMLGET is used for XML data, not general JSON-style object access.

Official Snowflake documentation reference:

Snowflake documentation describes GET_PATH as a function that extracts a value from semi-structured data using a path name.

Reference: Snowflake Documentation - GET_PATH; Querying semi-structured data; SnowPro Core Study Guide - Working with Semi-Structured Data.

NEW QUESTION: 395

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: ([SHOW ANSWER](#))

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: 396

Which Snowflake feature enables loading data from cloud storage as soon as files are available in a stage?

- A. Direct share
- B. Data replication
- C. copy_into < location > command
- D. Snowpipe

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 397

Which table function should be used to view details on a Directed Acyclic Graphic (DAG) run that is presently scheduled or is executing?

- A. TASK_HISTORY
- B. TASK_DEPENDENTS
- C. CURRENT_TASK_GRAPHS
- D. COMPLETE_TASK_GRAPHS

Answer: ([SHOW ANSWER](#))

The `CURRENT_TASK_GRAPH` stable function is designed to provide information on Directed Acyclic Graphs (DAGs) that are currently scheduled or executing within Snowflake. This function offers insights into the structure and status of task chains, enabling users to monitor and troubleshoot task executions. DAGs in Snowflake represent sequences of tasks with dependencies, and understanding their current state is crucial for managing complex workflows.

References: Snowflake Documentation on Task Management

NEW QUESTION: 398

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 399

Which key access control concept does Snowflake describe as a defined level of access to an object?

- A. Grant
- B. Privilege
- C. Role
- D. Session

Answer: ([SHOW ANSWER](#))

In Snowflake, the term "privilege" refers to a defined level of access to an object. Privileges are specific actions that roles can perform on securable objects in Snowflake, such as tables, views, warehouses, databases, and schemas. These privileges are granted to roles and can be further granted to users through their roles, forming the basis of Snowflake's access control framework.

References: Snowflake Documentation on Access Control Privileges

NEW QUESTION: 400

What metadata is stored for each micro-partition? (Select TWO)

- A. The average depth of values for each of the columns
- B. The clustering key used in the table
- C. The range of values for each of the columns in the full table
- D. The number of distinct values
- E. The minimum and maximum values for each of the columns

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 401

Which task is supported by the use of Access History in Snowflake?

- A. Data backups

- B. Cost monitoring
- C. Compliance auditing
- D. Performance optimization

Answer: (SHOW ANSWER)

Access History in Snowflake is primarily utilized for compliance auditing. The Access History feature provides detailed logs that track data access and modifications, including queries that read from or write to database objects. This information is crucial for organizations to meet regulatory requirements and to perform audits related to data access and usage.

Role of Access History: Access History logs are designed to help organizations understand who accessed what data and when. This is particularly important for compliance with various regulations that require detailed auditing capabilities.

How Access History Supports Compliance Auditing:

By providing a detailed log of access events, organizations can trace data access patterns, identify unauthorized access, and ensure that data handling complies with relevant data protection laws and regulations.

Access History can be queried to extract specific events, users, time frames, and accessed objects, making it an invaluable tool for compliance officers and auditors.

Reference: For more information on how Access History supports compliance auditing, refer to the Snowflake documentation on Access History: https://docs.snowflake.com/en/sql-reference/account-usage/access_history.html

NEW QUESTION: 402

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

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

Answer: (SHOW ANSWER)

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

References:

Snowflake Documentation on Clustering Keys & Clustered Tables¹.

Community discussions on how source data's ordering affects a table with a cluster key

NEW QUESTION: 403

Which Snowflake tool is recommended for data batch processing?

- A. SnowCD
- B. SnowSQL
- C. Snowsight
- D. The Snowflake API

Answer: (SHOW ANSWER)

For data batch processing in Snowflake, the recommended tool is:

B . SnowSQL: SnowSQL is the command-line client for Snowflake. It allows for executing SQL queries, scripts, and managing database objects. It's particularly suitable for batch processing tasks due to its ability to run SQL scripts that can execute multiple commands or queries in sequence, making it ideal for automated or scheduled tasks that require bulk data operations.

SnowSQL provides a flexible and powerful way to interact with Snowflake, supporting operations such as loading and unloading data, executing complex queries, and managing Snowflake objects from the command line or through scripts.

References:

Snowflake Documentation: SnowSQL (CLI Client) at Snowflake Documentation

NEW QUESTION: 404

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

Which condition will trigger a cluster to shut-down?

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

Answer: (SHOW ANSWER)

In a standard multi-cluster warehouse with auto-scaling, a cluster will shut down when, after 2-3 consecutive checks, the system determines that the load on the least-loaded cluster could be redistributed to other clusters.

This ensures efficient resource utilization and cost management. References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 405

What is the MINIMUM size of a table for which Snowflake recommends considering adding a clustering key?

- A. 1 Kilobyte (KB)
- B. 1 Megabyte (MB)
- C. 1 Gigabyte (GB)
- D. 1 Terabyte (TB)

Answer: (SHOW ANSWER)

Snowflake recommends considering adding a clustering key to a table when its size reaches 1 Terabyte (TB) or larger. Clustering keys help optimize the storage and query performance by organizing the data in a table based on the specified columns. This is particularly beneficial for large tables where data retrieval can become inefficient without proper clustering.

Why Clustering Keys Are Important: Clustering keys ensure that data stored in Snowflake is physically ordered in a way that aligns with the most frequent access patterns, thereby reducing the amount of scanned data during queries and improving performance.

Recommendation Basis: The recommendation for tables of size 1 TB or larger is based on the observation that smaller tables generally do not benefit as much from clustering, given Snowflake 's architecture. However, as tables grow in size, the benefits of clustering become more pronounced.

Implementing Clustering Keys:

To set a clustering key for a table, you can use the `CLUSTER BY` clause during table creation or alter an existing table to add it:

```
CREATE TABLE my_table (... ) CLUSTER BY (column1, column2);
```

Or for an existing table:

```
ALTER TABLE my_table CLUSTER BY (column1, column2);
```

Reference: For additional information on clustering in Snowflake and recommendations on when to use clustering keys, consult the Snowflake documentation on managing and optimizing data clustering:

<https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html>

NEW QUESTION: 406

Which privilege is required to grant a database role to another role?

- A. USAGE on the database
- B. OWNERSHIP on the database role
- C. CREATE ROLE on the account
- D. MANAGE GRANTS on the account

Answer: (SHOW ANSWER)

The correct answer is B. OWNERSHIP on the database role .

A database role is a securable object within a database. To grant a database role to another role, the granting role must have the authority to manage that database role. The OWNERSHIP privilege on the database role gives full control over the database role, including the ability to grant it.

Why B is correct:

The owner of a database role can grant that database role to another database role or account role, depending on the supported grant relationship.

Example concept:

```
GRANT DATABASE ROLE database1.db_role1 TO ROLE account_role1;
```

To perform this action, the executing role must have appropriate authority over database1.db_role1, such as OWNERSHIP.

Why the other options are incorrect:

- A). USAGE on the database allows access to the database namespace, but it does not grant authority to grant a database role.
- C). CREATE ROLE applies to creating account roles, not granting database roles.
- D). MANAGE GRANTS can allow broad grant management, but the specific required object-level privilege in this question is OWNERSHIP on the database role.

Official Snowflake documentation reference:

Snowflake documentation explains that database roles are securable objects and that ownership provides full control over the object, including grant management.

Reference: Snowflake Documentation - Database roles; Snowflake Documentation - GRANT DATABASE ROLE; Snowflake Documentation - Access control privileges; SnowPro Core Study Guide - Security and Access Control.

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here: <https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (**1187** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 407

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

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

Answer: (SHOW ANSWER)

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

NEW QUESTION: 408

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

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

Answer: (SHOW ANSWER)

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

NEW QUESTION: 409

What do temporary and transient tables have in common in Snowflake? (Select TWO).

- A. Both tables have no Fail-safe period.
- B. Both tables have data retention period maximums of one day.
- C. Both tables are visible only to a single user session.
- D. For both tables the retention period ends when the tables are dropped.

E. For both tables, the retention period does not end when the session ends

Answer: ([SHOW ANSWER](#))

Temporary and transient tables in Snowflake share several characteristics, notably, neither table type has a Fail-safe period. Fail-safe is a feature that provides additional data protection beyond the Time Travel period.

However, this feature does not apply to temporary or transient tables. Additionally, for both types of tables, the data retention period effectively ends when the tables are dropped. This means that once these tables are deleted, their data is not recoverable, distinguishing them from permanent tables, which benefit from Snowflake ' s Time Travel and Fail-safe features.

References: Snowflake Documentation on Table Types

NEW QUESTION: 410

Which objects will incur storage costs associated with Fail-safe?

- A. External tables
- B. Permanent tables
- C. Data files available in internal stages
- D. Data files available in external stages

Answer: ([SHOW ANSWER](#))

Snowflake ' s Data Exchange allows users to create and manage a group of accounts to which they can offer data shares. This platform facilitates secure and governed data sharing within and between organizations.

Data Exchange Setup:Users can set up a Data Exchange and invite other Snowflake accounts to join.

Sharing Data:Within the Data Exchange, users can offer shares to the entire group or to specific accounts.

Governance and Security:Data Exchange provides tools for data governance and security, ensuring that shared data is only accessible to authorized accounts.

References:

Snowflake Documentation: Snowflake Data Exchange

Snowflake Documentation: Managing a Data Exchange

NEW QUESTION: 411

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

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

Answer: ([SHOW ANSWER](#))

Snowflake's architecture is designed to operate entirely on cloud infrastructure. It uses the core compute and storage services of each cloud provider, which allows it to leverage the scalability and reliability of cloud resources. Snowflake's

services are deployed across multiple availability zones within a cloud provider's region to ensure high availability and fault tolerance. References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 412

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

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

Answer: ([SHOW ANSWER](#))

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

References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 413

How does Snowflake recommend defining a clustering key on a high-cardinality column that includes a 15 digit ID numbered column ID_NUMBER?

- A. TRUNC (ID_NUMBER, 5)
- B. TO CHAR(ID_NUMBER)
- C. ID_NUMBER*100
- D. TRUNC (ID_NUMBER, -6)

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 414

How can customer data be restricted by country for a global company using a multi-tenant environment?

- A. Implement a secure User-Defined Function (UDF) and use a secure data share for each country.
- B. Implement a column-level security policy in the customer table that restricts access based on the user 's geographic location.
- C. Implement a new account structure and create sub-accounts for each country. Then replicate the required data.
- D. Implement a row-level security policy in the customer table that looks at a mapping table identifying which users should have access to what data.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 415

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: (SHOW ANSWER)

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: 416

Which operator is used to evaluate query pruning efficiency by comparing partitions scanned to partitions total?

- A. Filter
- B. IndexScan
- C. ExternalScan
- D. TableScan

Answer: (SHOW ANSWER)

NEW QUESTION: 417

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

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

D. USAGE

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 418

How can a Snowflake user post-process the result of SHOW FILE FORMATS?

- A. Use the RESULT_SCAN function.
- B. Create a CURSOR for the command.
- C. Put it in the FROM clause in brackets.
- D. Assign the command to RESULTSET.

Answer: ([SHOW ANSWER](#))

first run SHOW FILE FORMATS

then SELECT * FROM TABLE(RESULT_SCAN(LAST_QUERY_ID(-1)))

https://docs.snowflake.com/en/sql-reference/functions/result_scan#usage-notes

NEW QUESTION: 419

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: ([SHOW ANSWER](#))

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: 420

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: ([SHOW ANSWER](#))

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: 421

Why would the REPEATABLE keyword be used in a SAMPLE clause?

- A. To limit the sample to a specific number of rows
- B. To ensure no duplicates are returned when a query is run
- C. To ensure the same sample is returned each time a query is run
- D. To force the sampling to be performed on blocks rather than rows

Answer: ([SHOW ANSWER](#))

The correct answer is C. To ensure the same sample is returned each time a query is run .

The REPEATABLE keyword is used with a seed value to make sampling deterministic. This means Snowflake can return the same sample when the query is rerun with the same seed and the underlying data has not changed.

Why C is correct:

```
SELECT *
```

```
FROM table_A
```

```
SAMPLE BERNOULLI (10) REPEATABLE (123);
```

This returns a repeatable sample based on the seed value 123.

Why the other options are incorrect:

- A). Returning a fixed number of rows uses SAMPLE (< n > ROWS), not REPEATABLE.
- B). REPEATABLE does not remove duplicates.
- D). Block sampling is controlled with BLOCK or SYSTEM, not REPEATABLE.

Official Snowflake documentation reference:

Snowflake documentation explains that REPEATABLE or SEED specifies a seed value to make sample results repeatable.

Reference: Snowflake Documentation - SAMPLE / TABLESAMPLE; SnowPro Core Study Guide - SQL and Snowflake Objects.

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 422

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 423

Authorization to execute CREATE < object > statements comes only from which role?

- A. Primary role
- B. Secondary role
- C. Application role
- D. Database role

Answer: ([SHOW ANSWER](#))

In Snowflake, the authorization to execute CREATE < object > statements, such as creating tables, views, databases, etc., is determined by the role currently set as the user 's primary role. The primary role of a user or session specifies the set of privileges (including creation privileges) that the user has. While users can have multiple roles, only the primary role is used to determine what objects the user can create unless explicitly specified in the session.

Reference: This is based on the principle of Role-Based Access Control (RBAC) in Snowflake, where roles are used to manage access permissions. The official Snowflake documentation on Understanding and Using Roles would be the best resource to verify this information: <https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#roles>

NEW QUESTION: 424

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: ([SHOW ANSWER](#))

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: 425

What is the default value in the Snowflake Web Interface (UI) for auto suspending a Virtual Warehouse?

- A. 1 minutes
- B. 5 minutes
- C. 10 minutes
- D. 15 minutes

Answer: C ([LEAVE A REPLY](#))

The default value for auto-suspending a Virtual Warehouse in the Snowflake Web Interface (UI) is 10 minutes. This setting helps manage compute costs by automatically suspending warehouses that are not in use, ensuring that compute resources are efficiently allocated and not wasted on idle warehouses.

References:

Snowflake Documentation: Virtual Warehouses

NEW QUESTION: 426

Which column data type is supported when defining a clustering key on a table?

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

Answer: ([SHOW ANSWER](#)**)**

The correct answer is A. VARCHAR .

A clustering key can be defined on supported table columns or expressions to help Snowflake organize micro-partitions and improve pruning for certain query patterns. Standard scalar data types such as VARCHAR are supported.

Why A is correct:

VARCHAR is a standard scalar data type and can be used in a clustering key.

Example:

```
CREATE TABLE customer_orders (  
  customer_id VARCHAR,  
  order_date DATE,  
  amount NUMBER  
)  
CLUSTER BY (customer_id);
```

Why the other options are incorrect:

B). VARIANT stores semi-structured data. Clustering directly on raw semi-structured columns is not the best answer here; clustering keys should use supported scalar columns or expressions.

C). OBJECT is a semi-structured data type and is not the correct choice for a clustering key column in this question.

D). GEOGRAPHY is not the correct supported clustering key data type in this context.

Official Snowflake documentation reference:

Snowflake documentation explains that clustering keys can be defined using table columns or expressions, and standard scalar columns such as string, number, and date/time columns are commonly used for clustering.

Reference: Snowflake Documentation - Clustering keys and clustered tables; Snowflake Documentation - Table data types; SnowPro Core Study Guide - Performance Optimization.

NEW QUESTION: 427

What information is stored in the ACCESS_HISTORY view?

- A. History of the files that have been loaded into Snowflake
- B. Names and owners of the roles that are currently enabled in the session
- C. Query details such as the objects included and the user who executed the query
- D. Details around the privileges that have been granted for all objects in an account

Answer: (SHOW ANSWER)

Details around the privileges that have been granted for all objects in an account.

The ACCESS_HISTORY view in Snowflake provides a comprehensive log of access control changes, including grants and revocations of privileges on all securable objects within the account. This information is crucial for auditing and monitoring the security posture of your Snowflake environment.

Here's how to understand and use the ACCESS_HISTORY view:

Purpose of ACCESS_HISTORY View: It is designed to track changes in access controls, such as when a user or role is granted or revoked privileges on various Snowflake objects. This includes tables, schemas, databases, and more.

Querying ACCESS_HISTORY: To access this view, you can use the following SQL query pattern:

```
SELECT * FROM SNOWFLAKE.ACCOUNT_USAGE.ACCESS_HISTORY WHERE EVENT_TYPE = 'GRANT'  
OR EVENT_TYPE = 'REVOKE';
```

Interpreting the Results: The results from the `ACCESS_HISTORY` view include the object type, the specific privilege granted or revoked, the grantee (who received or lost the privilege), and the timestamp of the event.

This data is invaluable for audits and compliance checks.

Reference: For detailed information on the `ACCESS_HISTORY` view and how to interpret its data, please visit the official Snowflake documentation: https://docs.snowflake.com/en/sql-reference/account-usage/access_history.html

NEW QUESTION: 428

A table was created using this statement:

```
CREATE TRANSIENT TABLE T1 (  
  ID NUMERIC NOT NULL,  
  NAME STRING COLLATE 'en-ci') WITH MASKING POLICY MP1  
)  
CLUSTER BY ID  
WITH ROW ACCESS POLICY RPI1 ON (NAME)
```

Which feature of this table will incur both compute and storage costs?

- A. The row access policy
- B. The Dynamic Data Masking policy
- C. The not null constraint
- D. The Automatic Clustering

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 429

What is the MINIMUM Snowflake edition that supports database replication?

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

Answer: ([SHOW ANSWER](#))

The minimum Snowflake edition that supports database replication is the Enterprise edition. Database replication allows data to be replicated between different Snowflake accounts or regions, providing high availability and disaster recovery capabilities.

Snowflake Documentation: Database Replication

NEW QUESTION: 430

How can a user work with multiple assigned roles without the need to frequently switch roles?

- A. Use secondary roles and primary roles.
- B. Grant the user the `ACCOUNTADMIN` role.
- C. Grant the user the parent role of the roles that have been assigned.
- D. Create a new role and grant it to the user, then grant the existing roles to the new role.

Answer: ([SHOW ANSWER](#))

The correct answer is A. Use secondary roles and primary roles .

Corrected typo:

Option B contains a typing error. ACCOONTADKIN should be corrected to ACCOUNTADMIN .

Snowflake allows a session to use a primary role and, optionally, secondary roles. Secondary roles allow users to access privileges from additional granted roles without constantly switching the active primary role.

Why A is correct:

A user can activate secondary roles by using USE SECONDARY ROLES. When secondary roles are enabled, the user can use privileges granted to all active secondary roles in addition to the current primary role.

Example:

```
USE ROLE analyst_role;
```

```
USE SECONDARY ROLES ALL;
```

This allows the user to operate with privileges from the primary role and all granted secondary roles.

Why the other options are incorrect:

B). Granting ACCOUNTADMIN is excessive and violates least privilege.

C). Granting a parent role may work in some role hierarchy designs, but it is not the specific Snowflake feature that allows multiple roles to be active in a session.

D). Creating a new combined role can work administratively, but it adds operational overhead and is not as direct as using secondary roles.

Official Snowflake documentation reference:

Snowflake documentation explains that a session has one current primary role and can also activate secondary roles.

Secondary roles allow the privileges of additional roles granted to the user to be considered during authorization.

Reference: Snowflake Documentation - Role hierarchy and privilege inheritance; Snowflake Documentation

- Secondary roles; SnowPro Core Study Guide - Security and Access Control.

NEW QUESTION: 431

What kind of authentication do Snowpipe REST endpoints use?

A. OAuth

B. Key-based

C. Username and password

D. Single Sign-On (SSO)

Answer: (SHOW ANSWER)

Snowpipe uses key-based authentication for its REST endpoints. This involves generating and using a key pair (public and private keys) to securely authenticate API requests.

Generate Key Pair: Generate a public and private key pair.

Register Public Key: Register the public key with the Snowflake user that will be making the API requests.

Authenticate Requests: Use the private key to sign API requests sent to Snowpipe REST endpoints.

References:

Snowflake Documentation: Key Pair Authentication & Key Rotation

Snowflake Documentation: Using Snowpipe REST API

NEW QUESTION: 432

Which of the following describes the Snowflake Cloud Services layer?

A. Coordinates activities in the Snowflake account

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 433

Which of the following are considerations when using a directory table when working with unstructured data?

(Choose two.)

- A. A directory table is a separate database object.
- B. Directory tables store data file metadata.
- C. A directory table will be automatically added to a stage.
- D. Directory tables do not have their own grantable privileges.
- E. Directory table data can not be refreshed manually.

Answer: ([SHOW ANSWER](#))

Directory tables in Snowflake are used to store metadata about data files in a stage. They are not separate database objects but are conceptually similar to external tables. Directory tables do not have grantable privileges of their own

NEW QUESTION: 434

What impact does dropping a column from a table have on the micro-partitions of that table?

- A. The existing micro-partitions will not be affected.
- B. New micro-partitions will be created containing only the data for the remaining columns.
- C. The data associated with the dropped column will be removed from all existing micro-partitions.
- D. The current micro-partitions will not be available for Time Travel, and new micro-partitions will be created.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 435

A user creates a stage using the following command:

What will be the outcome?

- A. An error will be received stating that the storage location for the stage must be identified when creating a stage with a directory table.
- B. The command will fail to run because the name of the directory table is not specified.
- C. A stage with a directory table that has metadata that must be manually refreshed will be created.
- D. A stage with a directory table set to automatically refresh will be created.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 436

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

- A. Into single files with 100-250 MB of compressed data per file

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

Answer: ([SHOW ANSWER](#))

When loading data into Snowflake, it is recommended to organize the data into single files with 100-250 MB of compressed data per file. This size range is optimal for parallel processing and can help in achieving better performance during data loading operations. References: [COF-C03] SnowPro Core Certification Exam Study Guide

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (**1187** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 437

What are the possible values within a METADATASACTION column in a Snowflake stream? (Select TWO).

- A. INSERT
- B. UPDATE
- C. DELETE
- D. TRUNCATE
- E. UPSERT

Answer: ([SHOW ANSWER](#))

In Snowflake streams, the METADATASACTION column indicates the type of data manipulation operation that has occurred. The possible values include INSERT and DELETE.

INSERT: Indicates that a new row has been inserted into the table.

DELETE: Indicates that a row has been deleted from the table.

References:

Snowflake Documentation: Change Data Capture (CDC) with Streams

Snowflake Documentation: Stream Data Capture

NEW QUESTION: 438

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 439

Who can create and manage reader accounts? (Select TWO).

- A. A user with ACCOUNTADMIN role
- B. A user with securityadmin role
- C. A user with SYSADMIN role
- D. A user with ORGADMIN role
- E. A user with CREATE ACCOUNT privilege

Answer: ([SHOW ANSWER](#))

In Snowflake, reader accounts are special types of accounts that allow data sharing with external consumers without them having their own Snowflake account. The creation and management of reader accounts can be performed by users with the ACCOUNTADMIN role or the ORGADMIN role. The ACCOUNTADMIN role has comprehensive administrative privileges within a Snowflake account, including managing other accounts and roles. The ORGADMIN role, which is higher in hierarchy, oversees multiple accounts within an organization and can manage reader accounts across those accounts.

Snowflake Documentation: Creating and Managing Reader Accounts

NEW QUESTION: 440

Which role is responsible for managing the billing and credit data within Snowflake?

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

Answer: ([SHOW ANSWER](#))

The ORGADMIN role in Snowflake is responsible for managing organization-level administrative functions, which include managing billing and credit data. This role has the highest level of administrative privileges and can oversee multiple Snowflake accounts within an organization.

Snowflake Documentation: Account and Organization Roles

NEW QUESTION: 441

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: ([SHOW ANSWER](#))

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 tables.

NEW QUESTION: 442

Which option will replace NULL values with an empty string while unloading data into a file?

- A. `null_if= ( ' ' )`
- B. `null_if=( ' NULL ' )`
- C. `EMPTY_FIELD_AS_NULL = TRUE`
- D. `FIELD_OPTIONALLY_ENCLOSED_BY = NONE`

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 443

What are characteristics of Secure Data Sharing? Select TWO.

- A. Allows regular views to be shared.
- B. Sensitive data is automatically masked.
- C. Data is not transferred between accounts.
- D. Consumer account will incur compute costs.
- E. Source data can be edited by users with reader accounts.

Answer: ([SHOW ANSWER](#))

The correct answers are C. Data is not transferred between accounts and D. Consumer account will incur compute costs .

Important correction:

The provided answer A, B is not correct.

Secure Data Sharing allows data providers to share selected database objects with consumer accounts without copying or transferring the underlying data. Consumers query the shared data using their own virtual warehouses, so the consumer generally pays for the compute used to query the share.

Why C is correct:

With Secure Data Sharing, the provider's data remains in the provider account. Snowflake shares metadata and access, not a physical copy of the data.

Why D is correct:

Consumers use their own compute resources to query shared data. Therefore, the consumer account incurs compute costs when querying the share.

Why the other options are incorrect:

- A). Regular, non-secure views generally cannot be shared in the standard secure sharing model. Views shared with consumers should be secure views, unless specific account-level settings allow non-secure objects. The general SnowPro Core concept is that secure views are used to share view logic safely.
- B). Sensitive data is not automatically masked simply because data is shared. Masking policies or secure views must be designed and applied intentionally.
- E). Reader accounts can query shared data but cannot edit the provider's source data.

Official Snowflake documentation reference:

Snowflake documentation explains that Secure Data Sharing does not copy or move data between accounts. Consumers use their own warehouses to query shared data, while providers retain control of the shared objects. Reference: Snowflake Documentation - Secure Data Sharing; Snowflake Documentation - Shares and consumer accounts; Snowflake Documentation - Secure views; SnowPro Core Study Guide - Data Protection and Governance.

NEW QUESTION: 444

When cloning tables, which INFORMATION_SCHEMA view will show different columns for owned storage and referenced storage?

- A. DATABASE_STORAGE_USAGE_HISTORY
- B. STORAGE_USAGE
- C. TABLE_STORAGE_METRICS
- D. STAGE_STORAGE_USAGE_HISTORY

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 445

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 446

What does the LATERAL modifier for the FLATTEN function do?

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

Answer: ([SHOW ANSWER](#))

The LATERAL modifier for the FLATTEN function allows joining information outside the object (such as other columns in the source table) with the flattened data, creating a lateral view that correlates with the preceding tables in the FROM clause²³⁴⁵. References: [COF-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 447

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

- A. ORGADMIN
- B. USERADMIN

C. SYSADMIN

D. SECURITYADMIN

Answer: ([SHOW ANSWER](#))

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: 448

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

A. Data can be filtered by an optional where clause

B. Incoming data can be joined with other tables

C. Columns can be reordered

D. Columns can be omitted

E. Row level access can be defined

Answer: ([SHOW ANSWER](#))

In a CREATE PIPE ... AS COPY ... FROM (....) statement, the supported transformations include filtering data using an optional WHERE clause and omitting columns. The WHERE clause allows for the specification of conditions to filter the data that is being loaded, ensuring only relevant data is inserted into the table.

Omitting columns enables the exclusion of certain columns from the data load, which can be useful when the incoming data contains more columns than are needed for the target table.

References:

[COF-C03] SnowPro Core Certification Exam Study Guide

Simple Transformations During a Load¹

NEW QUESTION: 449

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

A. Create a zero-copy clone of the worksheet and grant permissions to teammates

B. Create a private Data Exchange so that any teammate can use the worksheet

C. Share the worksheet with teammates within Snowsight

D. Create a database and grant all permissions to teammates

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 450

A user executes a long-running query at the same time every morning. A second user with the same role needs to access the query results for additional data processing and analysis. How should the results be reproduced for the second user?

A. Locate the query ID from QUERY_HISTORY and get the results using the RESULT_SCAN function.

B. Locate the query text from QUERY_HISTORY and rerun the query to produce the results.

C. Get the results using the RESULT_SCAN function with -1.

D. Locate the task that executes the query and rerun the query to produce the results.

Answer: ([SHOW ANSWER](#))

RESULT_SCAN can return the result set of a previous query when the query ID is supplied, provided the result is still available and the user has appropriate access. Since the second user has the same role, locating the query ID and using RESULT_SCAN is the most efficient way to reproduce the results without rerunning the long query.

Example:

```
SELECT *  
FROM TABLE(RESULT_SCAN( ' < query_id > ' ));
```

Reference: Snowflake Documentation - RESULT_SCAN; Persisted query results; QUERY_HISTORY.

NEW QUESTION: 451

Which of the following are characteristics of security in Snowflake?

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

Answer: ([SHOW ANSWER](#))

One of the security features of Snowflake includes the periodic rekeying of encrypted data, which is available with the Snowflake Enterprise edition and higher². This ensures that the encryption keys are rotated regularly to maintain a high level of security. References: [COF-C03] SnowPro Core Certification Exam Study Guide

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (**1187** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 452

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

- A.** True
- B.** False

Answer: ([SHOW ANSWER](#))

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

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

NEW QUESTION: 453

How can data be shared in Snowflake with another account located in the same region without the need to copy or move the data?

- A. Use the Snowflake Marketplace
- B. Use a stream
- C. Use a Data Exchange
- D. Use a direct share

Answer: ([SHOW ANSWER](#))

Direct Shares allow secure data sharing between Snowflake accounts in the same region . No data is copied; Snowflake provides access to the original data using metadata and access control.

Reference:

Snowflake Docs: Data Sharing

NEW QUESTION: 454

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 455

Which functions are used to manipulate semi-structured data? (Select TWO)

- A. WIDTH_BUCKET
- B. HASH_AGG
- C. ARRAY_APPEND
- D. OBJECT_INSERT
- E. LISTAGG

Answer: ([SHOW ANSWER](#))

Functions like ARRAY_APPEND() and OBJECT_INSERT() are system-defined for manipulating ARRAY and OBJECT types within VARIANT data structures.

Reference:

Snowflake Docs: Functions for Semi-Structured Data

NEW QUESTION: 456

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

- A. grant USAGE on all tables in schema DB1.SCHEMA to role MYROLE;

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

Answer: C,D ([LEAVE A REPLY](#))

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

NEW QUESTION: 457

Which command will provide an account with access to a share?

- A. ALTER SHARE
- B. CREATE SHARE
- C. DESCRIBE SHARE
- D. GRANT < privilege > ... TO SHARE

Answer: ([SHOW ANSWER](#))

The correct answer is A. ALTER SHARE .

To provide a consumer account with access to an existing share, the provider adds the consumer account to the share. This is done using the ALTER SHARE command with the ADD ACCOUNTS clause.

Example:

```
ALTER SHARE my_share  
ADD ACCOUNTS = consumer_org.consumer_account;
```

Why A is correct:

ALTER SHARE modifies an existing share. Adding consumer accounts to a share gives those accounts access to the share.

Why the other options are incorrect:

- B). CREATE SHARE creates a new share object, but it does not by itself add a consumer account to an existing share.
- C). DESCRIBE SHARE displays metadata about a share; it does not grant access.
- D). GRANT < privilege > ... TO SHARE grants database object privileges to the share, such as usage on a database or select on a table. It does not add a consumer account to the share.

Official Snowflake documentation reference:

Snowflake documentation explains that ALTER SHARE ... ADD ACCOUNTS is used to add consumer accounts to a share. Object privileges are granted separately to the share using GRANT ... TO SHARE.

Reference: Snowflake Documentation - ALTER SHARE; Snowflake Documentation - Secure Data Sharing; Snowflake Documentation - GRANT ... TO SHARE; SnowPro Core Study Guide - Data Protection and Governance.

NEW QUESTION: 458

What can be used to process unstructured data?

- A. External tables
- B. The copy into < table > command

C. External functions

D. Snowpipe

Answer: ([SHOW ANSWER](#))

To process unstructured data in Snowflake, external functions can be used.

External Functions: These allow you to call external services and processing engines from within Snowflake SQL.

External functions can be used to handle complex processing tasks that are not natively supported by Snowflake, including those involving unstructured data.

Implementation: You define an external function in Snowflake that points to an external processing service (e.g., AWS Lambda, Google Cloud Functions).

Snowflake Documentation on External Functions

NEW QUESTION: 459

If a table was dropped and the DATA_RETENTION_TIME_IN_DAYS was set to 0, how can the data from the table be recovered?

A. By creating a failover group

B. By creating a Snowflake Support ticket

C. By running the UNDROP TABLE command

D. By running the CLONE ... BEFORE command

Answer: B ([LEAVE A REPLY](#))

The correct answer is B. By creating a Snowflake Support ticket .

If DATA_RETENTION_TIME_IN_DAYS = 0, Time Travel retention is disabled for the table. This means the table cannot be recovered using normal Time Travel operations such as UNDROP TABLE or CLONE ... BEFORE.

Why B is correct:

After Time Travel retention is unavailable, recovery may only be possible through Snowflake Fail-safe for permanent tables. Fail-safe is not directly accessible by users. Recovery from Fail-safe requires contacting Snowflake Support.

Important qualification:

This assumes the dropped table is a permanent table and the data is still within Fail-safe. Temporary and transient tables do not have Fail-safe, so they cannot be recovered this way.

Why the other options are incorrect:

A). A failover group is used for replication and disaster recovery across accounts or regions. It does not recover a dropped table with no Time Travel retention.

C). UNDROP TABLE relies on Time Travel retention. With retention set to 0, this is not available.

D). CLONE ... BEFORE also relies on Time Travel. It cannot recover data when no retention period exists.

Official Snowflake documentation reference:

Snowflake documentation explains that Time Travel enables users to restore dropped objects within the retention period. Fail-safe is a separate recovery mechanism for permanent tables and is only available through Snowflake Support.

Reference: Snowflake Documentation - Time Travel; Fail-safe; UNDROP TABLE; SnowPro Core Study Guide - Data Protection and Governance.

NEW QUESTION: 460

When should a multi-cluster virtual warehouse be used in Snowflake?

- A. When queuing is delaying query execution on the warehouse
- B. When there is significant disk spilling shown on the Query Profile
- C. When dynamic vertical scaling is being used in the warehouse
- D. When there are no concurrent queries running on the warehouse

Answer: ([SHOW ANSWER](#))

A multi-cluster virtual warehouse in Snowflake is designed to handle high concurrency and workload demands by allowing multiple clusters of compute resources to operate simultaneously. The correct scenario to use a multi-cluster virtual warehouse is:

A . When queuing is delaying query execution on the warehouse: Multi-cluster warehouses are ideal when the demand for compute resources exceeds the capacity of a single cluster, leading to query queuing. By enabling additional clusters, you can distribute the workload across multiple compute clusters, thereby reducing queuing and improving query performance.

This is especially useful in scenarios with fluctuating workloads or where it's critical to maintain low response times for a large number of concurrent queries.

References:

Snowflake Documentation: Multi-Cluster Warehouses at [Snowflake Documentation](#)

NEW QUESTION: 461

When referring to User-Defined Function (UDF) names in Snowflake, what does the term overloading mean?

- A. There are multiple SQL UDFs with the same names and the same number of arguments.
- B. There are multiple SQL UDFs with the same names and the same number of argument types.
- C. There are multiple SQL UDFs with the same names but with a different number of arguments or argument types.
- D. There are multiple SQL UDFs with different names but the same number of arguments or argument types.

Answer: ([SHOW ANSWER](#))

In Snowflake, overloading refers to the creation of multiple User-Defined Functions (UDFs) with the same name but differing in the number or types of their arguments. This feature allows for more flexible function usage, as Snowflake can differentiate between functions based on the context of their invocation, such as the types or the number of arguments passed. Overloading helps to create more adaptable and readable code, as the same function name can be used for similar operations on different types of data.

References:

Snowflake Documentation: User-Defined Functions

NEW QUESTION: 462

Which role can create a database role?

- A. Any user granted the SECURITYADMIN role
- B. A custom role granted to the SECURITYADMIN role
- C. A role with the OWNERSHIP privilege on the database
- D. Any user granted the SYSADMIN role

Answer: ([SHOW ANSWER](#))

The correct answer is C. A role with the OWNERSHIP privilege on the database .

A database role is scoped to a specific database. To create a database role, a role must have sufficient privileges on the database where the database role will be created.

Why C is correct:

The role that owns the database has control over that database and can create database roles within it.

Example:

```
CREATE DATABASE ROLE my_database.my_db_role;
```

To run this successfully, the executing role must have the required privilege on the database, such as OWNERSHIP.

Why the other options are incorrect:

- A). SECURITYADMIN manages users, roles, and grants at the account level, but database role creation depends on privileges on the specific database.
- B). A custom role being granted to SECURITYADMIN does not automatically mean it can create database roles in a database.
- D). SYSADMIN commonly owns many account objects by default in recommended role models, but the required condition is privilege on the database, not simply holding the SYSADMIN role.

Official Snowflake documentation reference:

Snowflake documentation explains that database roles are created inside databases and require appropriate privileges on the database.

Reference: Snowflake Documentation - Database roles; CREATE DATABASE ROLE; Access control privileges; SnowPro Core Study Guide - Security and Access Control.

NEW QUESTION: 463

Which command should be used to drop files from an internal or external stage?

- A. DELETE
- B. TRUNCATE
- C. DROP
- D. REMOVE

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 464

What happens to deleted or modified records during the Time Travel retention period?

- A. They are moved to a separate table.
- B. They are permanently deleted.
- C. They are archived in a backup.
- D. They are versioned for recovery.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 465

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: ([SHOW ANSWER](#))

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: 466

Which Snowflake object does not consume any storage costs?

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

Answer: C ([LEAVE A REPLY](#))

Temporary tables do not consume any storage costs in Snowflake because they only exist for the duration of the session that created them and are automatically dropped when the session ends, thus incurring no long-term storage charges. References: [COF-C03] SnowPro Core Certification Exam Study Guide

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here:
<https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (1187 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 467

What is the most efficient method to share a subset of data from a table with a consumer account?

- A. Create a secure User-Defined Function, or UDF.
- B. Use a dynamic table.
- C. Create an external table.
- D. Create a secure view.

Answer: ([SHOW ANSWER](#))

The correct answer is D. Create a secure view .

A secure view is the recommended and efficient way to share a filtered subset of data with a consumer account using Snowflake Secure Data Sharing. The provider can define the rows and columns that should be visible, then share the secure view instead of sharing the full base table.

Why D is correct:

A secure view protects the view definition from unauthorized users and can be used to expose only selected rows or columns from an underlying table.

Example concept:

```
CREATE SECURE VIEW shared_sales_west AS
SELECT order_id, customer_id, amount, region
FROM sales
WHERE region = ' WEST ' ;
```

The provider can then include this secure view in a share, allowing the consumer to query only the approved subset of data.

Why the other options are incorrect:

- A). A secure UDF can protect function logic, but it is not the most efficient method for sharing a subset of table data.
- B). A dynamic table materializes transformation results, but it is not the standard method for secure data sharing with filtered access.
- C). An external table allows querying data in external cloud storage. It is not the best method for sharing a subset of Snowflake table data.

Official Snowflake documentation reference:

Snowflake documentation explains that secure views are used with Secure Data Sharing when a provider wants to share filtered or restricted data while protecting the underlying query definition.

Reference: Snowflake Documentation - Secure Data Sharing; Snowflake Documentation - Secure views; SnowPro Core Study Guide - Data Protection and Governance.

NEW QUESTION: 468

What is the difference between a stored procedure and a User-Defined Function (UDF)?

- A. Stored procedures can execute database operations while UDFs cannot.
- B. Returning a value is required in a stored procedure while returning values in a UDF is optional.
- C. Values returned by a stored procedure can be used directly in a SQL statement while the values returned by a UDF cannot.
- D. Multiple stored procedures can be called as part of a single executable statement while a single SQL statement can only call one UDF at a time.

Answer: (SHOW ANSWER)

Stored procedures in Snowflake can perform a variety of database operations, including DDL and DML, whereas UDFs are designed to return values and cannot execute database operations¹.

NEW QUESTION: 469

What is the purpose of a Query Profile?

- A. To profile how many times a particular query was executed and analyze its usage statistics over time.
- B. To profile a particular query to understand the mechanics of the query, its behavior, and performance.

C. To profile the user and/or executing role of a query and all privileges and policies applied on the objects within the query.

D. To profile which queries are running in each warehouse and identify proper warehouse utilization and sizing for better performance and cost balancing.

Answer: ([SHOW ANSWER](#))

The purpose of a Query Profile is to provide a detailed analysis of a particular query's execution plan, including the mechanics, behavior, and performance. It helps in identifying potential performance bottlenecks and areas for optimization

NEW QUESTION: 470

What is the purpose of the use of the VALIDATE command?

A. To view any queries that encountered an error

B. To verify that a SELECT query will run without error

C. To prevent a put statement from running if an error occurs

D. To see all errors from a previously run COPY INTO <table> statement

Answer: ([SHOW ANSWER](#))

The VALIDATE command in Snowflake is used to check for errors that occurred during the execution of a COPY INTO <table> statement. This command helps users identify and resolve data loading issues.

Run the COPY INTO Statement: Execute the COPY INTO <table> command to load data from a stage into a table.

```
COPY INTO my_table
```

```
FROM @my_stage
```

```
FILE_FORMAT = (FORMAT_NAME = 'my_format');
```

Validate the Load: Use the VALIDATE function to see if there were any errors during the data load.

```
SELECT *
```

```
FROM TABLE(VALIDATE(my_table, JOB_ID => 'my_copy_job_id'));
```

Review Errors: The VALIDATE function will return details about any errors that occurred, such as parsing errors or data type mismatches.

References:

Snowflake Documentation: Validating Data Loads

Snowflake Documentation: COPY INTO <table>

NEW QUESTION: 471

What should be considered before implementing a Data Clean Room?

A. SQL queries can be run against the data in the clean room.

B. Providers can only configure and share clean rooms with users in other Snowflake accounts.

C. Consumers must be approved by a clean room administrator before they can be invited.

D. All participating accounts must be in the same geographic region.

Answer: ([SHOW ANSWER](#))

A key consideration for Data Clean Rooms is account and region compatibility. Participating accounts generally need to be in supported regions for the clean room configuration and collaboration model.

Important note: Data Clean Rooms have evolved quickly, and regional/cross-cloud support can change. I cannot live-verify the newest Snowflake restrictions here, but this is the best SnowPro-style answer from the provided options. Reference: Snowflake Documentation - Snowflake Data Clean Rooms; Clean room limitations and considerations.

NEW QUESTION: 472

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

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

Answer: ([SHOW ANSWER](#))

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

References:

Continuous Data Protection | Snowflake Documentation¹

Data Storage Considerations | Snowflake Documentation²

Snowflake SnowPro Core Certification Study Guide³

Snowflake Data Cloud Glossary

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

NEW QUESTION: 473

Which statement will trigger a stream to advance its offset?

- A. DROP STREAM
- B. CREATE OR REPLACE STREAM
- C. DESCRIBE STREAM
- D. ALTER STREAM

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 474

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: ([SHOW ANSWER](#))

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: 475

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

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

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 476

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: **A,C** ([LEAVE A REPLY](#))

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-C03] SnowPro Core Certification Exam Study Guide

NEW QUESTION: 477

What Snowflake database object is derived from a query specification, stored for later use, and can speed up expensive aggregation on large data sets?

- A.** Temporary table
- B.** External table
- C.** Secure view
- D.** Materialized view

Answer: ([SHOW ANSWER](#))

A materialized view in Snowflake is a database object derived from a query specification, stored for later use, and can significantly speed up expensive aggregations on large data sets. Materialized views store the result of their underlying

query, reducing the need to recompute the result each time the view is accessed. This makes them ideal for improving the performance of read-heavy, aggregate-intensive queries.

References:

Snowflake Documentation: Using Materialized Views

NEW QUESTION: 478

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: ([SHOW ANSWER](#))

A cached MFA token is valid for up to four hours.<https://docs.snowflake.com/en/user-guide/security-mfa#using-mfa-token-caching-to-minimize-the-number-of-prompts-during-authentication-optional>

Valid COF-C03 Dumps shared by EduDump.com for Helping Passing COF-C03 Exam! EduDump.com now offer the **newest COF-C03 exam dumps**, the EduDump.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com COF-C03 dumps with Test Engine here: <https://www.edudump.com/exams/Snowflake/COF-C03/premium/> (**1187** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)