# Snowflake.COF-C03.v2026-11-21.q547

| | |
|---|---|
| **Exam Code:** | COF-C03 |
| **Exam Name:** | SnowPro® Core Certification 2026 Exam (COF-C03) |
| **Certification Provider:** | Snowflake |
| **Free Question Number:** | 547 |
| **Version:** | v2026-11-21 |
| **# of views:** | 110 |
| **# of Questions views:** | 6456 |
| https://www.freecram.net/torrent/Snowflake.COF-C03.v2026-11-21.q547.html | |

**NEW QUESTION: 1**
Which privilege allows a user to change the state of a virtual warehouse that was set to AUTO_RESUME = FALSE?
**A.** USAGE
**B.** APPLYBUDGET
**C.** OPERATE
**D.** MONITOR
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 2**
Which Snowflake feature provides increased login security for users connecting to Snowflake that is powered by Duo Security service?
**A.** OAuth
**B.** Network policies
**C.** Single Sign-On (SSO)
**D.** Multi-Factor Authentication (MFA)
**Answer: (SHOW ANSWER)**
Multi-Factor Authentication (MFA) provides increased login security for users connecting to Snowflake.
Snowflake's MFA is powered by Duo Security service, which adds an additional layer of security during the login process.

**NEW QUESTION: 3**
Which COPY INTO < table > parameter setting should be used to prevent a partial data load when an error is encountered?
**A.** FORCE = FALSE
**B.** RETURN_FAILED_ONLY = FALSE
**C.** LOAD_UNCERTAIN_FILES = FALSE
**D.** ON_ERROR = ABORT_STATEMENT

**Answer: (SHOW ANSWER)**
ON_ERROR = ABORT_STATEMENT instructs Snowflake to abort the COPY INTO statement when a data- loading error is encountered. This setting is used when the load should not continue after discovering an invalid record or file. Other ON_ERROR values, such as CONTINUE, can allow valid rows or files to load while problematic records are skipped, which can produce a partial data load.
FORCE controls whether files that were previously loaded are loaded again. LOAD_UNCERTAIN_FILES controls how files with uncertain load status are handled. RETURN_FAILED_ONLY controls result reporting and does not determine whether a load is atomic.
References:
Snowflake SQL Reference - COPY INTO < table > : https://docs.snowflake.com/en/sql-reference/sql/copy- into-table
Snowflake Documentation - Handling data-load errors: https://docs.snowflake.com/en/user-guide/data-load-considerations-load
==

**NEW QUESTION: 4**
What is the recommended file sizing for data loading using Snowpipe?
**A.** A compressed file size greater than 100 MB, and up to 250 MB
**B.** A compressed file size greater than 100 GB, and up to 250 GB
**C.** A compressed file size greater than 10 MB, and up to 100 MB
**D.** A compressed file size greater than 1 GB, and up to 2 GB
**Answer: (SHOW ANSWER)**
For data loading using Snowpipe, the recommended file size is a compressed file greater than 10 MB and up to 100 MB. This size range is optimal for Snowpipe's continuous, micro-batch loadingprocess, allowing for efficient and timely data ingestion without overwhelming the system with files that are too large or too small.
References:
[COF-C03] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Snowpipe1

**NEW QUESTION: 5**
Which command will identify the names of files that are in a Snowflake stage?
**A.** SELECT $1 FROM < location >
**B.** DESCRIBE FILE FORMAT
**C.** LIST
**D.** SHOW FILE FORMAT
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 6**
Files have been uploaded to a Snowflake internal stage. The files now need to be deleted.
Which SQL command should be used to delete the files?
**A.** PURGE
**B.** MODIFY

**C.** REMOVE
**D.** DELETE
**Answer: (SHOW ANSWER)**
The SQL command used to delete files from a Snowflake internal stage is REMOVE. This command can be used to remove files from either an internal or external stage within Snowflake

**NEW QUESTION: 7**
Which command is used to lake away staged files from a Snowflake stage after a successful data ingestion?
**A.** DELETE
**B.** DROP
**C.** REMOVE
**D.** TRUNCATE
**Answer: (SHOW ANSWER)**
TheREMOVEcommand 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.
UnlikeDELETE,DROP, orTRUNCATEcommands, which are used for managing data within Snowflake tables or dropping objects,REMOVEspecifically targets the management of files in stages.
References: Snowflake Documentation on Stages and File Management

**NEW QUESTION: 8**
Use of which file function allows a user to share unstructured data from an internal stage with an external reporting tool that does not have access to Snowflake " >
**A.** BUILD_SCOPED_FILE_URL
**B.** GET_PRESIGNED_URL
**C.** BUILD_STAGE_FILE_URL
**D.** GET_STAGE_LOCATION
**Answer: (SHOW ANSWER)**
TheGET_PRESIGNED_URLfunction in Snowflake generates a pre-signed URL for a file in an internal stage.
This URL can be shared with external tools or users who do not have direct access to Snowflake, allowing them to download the file.
Generate Pre-Signed URL:
SELECT GET_PRESIGNED_URL(@my_stage/file.txt);
Share the URL:The generated URL can be shared with external users or applications, enabling them to access the file directly.
References:
Snowflake Documentation: GET_PRESIGNED_URL
Snowflake Documentation: Working with Stages

**NEW QUESTION: 9**
What is the maximum total Continuous Data Protection (CDP) charges incurred for a temporary table?

**A.** 30 days
**B.** 7 days
**C.** 48 hours
**D.** 24 hours
**Answer: D (LEAVE A REPLY)**
For a temporary table, the maximum total Continuous Data Protection (CDP) charges incurred are forthe duration of the session in which the table was created, which does not exceed 24 hours2.
References = [ COF-C03 ] SnowPro Core Certification Exam Study Guide, Snowflake Documentation2

**NEW QUESTION: 10**
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.
Freecram.net
**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: 11**
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: (SHOW ANSWER)**
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: 12**
Which privilege is required for a role to be able to resume a suspended warehouse if auto-resume is not enabled?
**A.** USAGE
**B.** OPERATE
**C.** MONITOR
**D.** MODIFY
**Answer: (SHOW ANSWER)**
The OPERATE privilege is required for a role to resume a suspended warehouse if auto-resume is not enabled. This privilege allows the role to start, stop, suspend, or resume a virtual warehouse3.
References: [ COF-C03 ] SnowPro Core Certification E xam Study Guide

**NEW QUESTION: 13**
What step can reduce data spilling in Snowflake?
**A.** Using a larger virtual warehouse
**B.** Increasing the virtual warehouse maximum timeout limit
**C.** Increasing the amount of remote storage for the virtual warehouse
**D.** Using a common table expression (CTE) instead of a temporary table
**Answer: (SHOW ANSWER)**
To reduce data spilling in Snowflake, using a larger virtual warehouse is effective because it provides more memory and local disk space, whichcan accommodate larger data operations and minimize the need to spill data to disk or remote storage1. References: [COF-C03] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 14**
Which Snowflake object type can be used to restrict users from viewing the DDL associated with an object?
**A.** Secure view
**B.** Data masking policy
**C.** Stored procedure
**D.** Materialized view
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 15**
How many days is load history for Snowpipe retained?
**A.** 1 day
**B.** 7 days
**C.** 14 days
**D.** 64 days
**Answer: (SHOW ANSWER)**
Snowpipe retains load history for 14 days.This allows users to view and audit the data that has been loaded into Snowflake using Snowpipe within this time frame3.

**NEW QUESTION: 16**
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

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

**NEW QUESTION: 17**
Which Snowflake objects can be shared with other Snowflake accounts? (Choose three.)
**A.** Schemas
**B.** Roles
**C.** Secure Views
**D.** Stored Procedures
**E.** Tables
**F.** Secure User-Defined Functions (UDFs)
**Answer: (SHOW ANSWER)**
In Snowflake, you can share several types of objects with other Snowflake accounts. These include schemas, secure views, and secure user-defined functions (UDFs).Sharing these objects allows for collaboration and data access across different Snowflake accounts while maintaining security and governance controls4.

**NEW QUESTION: 18**
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 ofdata scanned, thereby improving query performance

**NEW QUESTION: 19**

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 setwill inherit the new retention period from the schema4.

**NEW QUESTION: 20**

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: 21**

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: 22**
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

**NEW QUESTION: 23**
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: 24**
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, theOPERATEprivilege is required on the pipe object. TheOPERATEprivilege allows users to perform operational tasks on specific objects such as pipes, tasks, and streams. Specifically, for a pipe, theOPERATEprivilege enables the user to execute theALTER PIPE ... SET PIPE_EXECUTION_PAUSED=TRUEorALTER PIPE ... SET PIPE_EXECUTION_PAUSED=FALSEcommands, 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:TheOPERATEprivilege is essential for pausing and resuming pipes. This privilege is more specific than general object access privileges likeSELECTorUSAGEand is tailored for operational control.
Granting the OPERATE Privilege:To grant theOPERATEprivilege on a pipe, an administrator or a user with the necessary grants can execute the SQL command:
GRANTOPERATEONPIPE < pipe_name > TOROLE < role_name > ;
Pausing and Resuming Pipes:Once theOPERATEprivilege 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: 25**
Which function can be used to convert semi-structured data into a relational representation?
**A.** ARRAYS_TO_OBJECT
**B.** OBJECT_KEYS
**C.** PARSE_JSON
**D.** FLATTEN
**Answer: D (LEAVE A REPLY)**

**NEW QUESTION: 26**
Which Snowflake feature enables loading data from cloud storage as soon as files ate available in a stage?
**A.** Snowpipe
**B.** Data replication
**C.** Direct share
**D.** COPY_INTO <location> command
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 27**
In addition to performing all the standard steps to share data, which privilege must be granted on each database referenced by a secure view in order to be shared?

**A.** READ
**B.** REFERENCES
**C.** REFERENCEJJSAGE
**D.** USAGE
**Answer: (SHOW ANSWER)**
In addition to performing all the standard steps to share data, theUSAGEprivilege must be granted on each database referenced by a secure view in order to be shared. When sharing a database or specific objects like secure views, the receiving account needs to have theUSAGEprivilege on the database and schema to access the shared data. Thisprivilege enables the receiving account to access the database and its schemas but does not allow for any DML operations. It ' s a prerequisite for accessing any objects within the database.
For a secure view to be part of a share, not only does the view itself need to be shared, but the underlying database (and schema, if applicable) must also be accessible to the recipients. GrantingUSAGEprivilege on the database ensures that the receiving account can access the database in a read-only mode to utilize the shared view.
References:
Snowflake Documentation on Shares: Creating and Managing Shares

**NEW QUESTION: 28**
Which step will reduce costs when backing up a table?
**A.** Clone the table.
**B.** Download the table data to a stage.
**C.** Share the table data with an account in another region.
**D.** Copy the table using a CREATE TABLE AS SELECT, or CTAS, statement.
**Answer: (SHOW ANSWER)**
The correct answer is A. Clone the table .
The most cost-efficient way to back up a Snowflake table is to use zero-copy cloning . A clone creates a copy of the table's metadata without immediately duplicating the underlying micro-partitions.
Why A is correct:
When a table is cloned, Snowflake creates a new independent table object that initially shares the same micro-partitions as the source table. Because the underlying data is not copied at clone creation time, cloning is fast and storage-efficient.
This is why cloning is commonly used for backups, development, testing, and point-in-time snapshots.
Important concept:
At the instant a clone is created, the original table and cloned table share the same micro-partitions. The clone becomes independent after creation, and changes to either object do not modify the other object. Additional storage is used only when data changes occur.
Why the other options are incorrect:
B). Downloading table data to a stage creates additional data files and increases storage and operational overhead.
C). Sharing data with another account in another region is not the simplest or lowest-cost table backup method. Cross-region replication or sharing can involve additional costs.
D). A CTAS statement physically creates a new table from query results and can duplicate data storage, making it less storage-efficient than cloning.

Official Snowflake documentation reference:
Snowflake documentation explains that zero-copy cloning creates a copy of a database, schema, or table without duplicating the underlying storage at the time the clone is created.
Reference: Snowflake Documentation - Zero-copy cloning; Snowflake Documentation - CREATE TABLE ... CLONE; SnowPro Core Study Guide - Snowflake Architecture and Data Protection.
==

**NEW QUESTION: 29**
When using SnowSQL, which configuration options are required when unloading data from a SQL query run on a local machine? {Select TWO).
**A.** echo
**B.** quiet
**C.** output_file
**D.** output_format
**E.** force_put_overwrite
**Answer: (SHOW ANSWER)**
When unloading data from SnowSQL (Snowflake ' s command-line client), to a file on a local machine, you need to specify certain configuration options to determine how and where the data should be outputted. The correct configuration options required are:
C). output_file: This configuration option specifies the file path where the output from the query should be stored. It is essential for directing the results of your SQL query into a local file, rather than just displaying it on the screen.
D). output_format: This option determines the format of the output file (e.g., CSV, JSON, etc.). It is crucial for ensuring that the data is unloaded in a structured format that meets the requirements of downstream processes or systems.
These options are specified in the SnowSQL configuration file or directly in the SnowSQL command line.
The configuration file allows users to set defaults and customize their usage of SnowSQL, including output preferences for unloading data.
References:
Snowflake Documentation: SnowSQL (CLI Client) at Snowflake Documentation Snowflake Documentation: Configuring SnowSQL at Snowflake Documentation

**NEW QUESTION: 30**
A team loads 50 TB of historical transaction data into a large table that continues to grow. A Snowflake practitioner runs a query that filters the table to a narrow date range:
WHERE transaction_date BETWEEN ' 2025-01-01 ' AND ' 2025-01-31 '
How do micro-partitions and their metadata improve this query's performance?
Freecram.net
**A.** Micro-partitions must be manually defined by applying a PARTITION BY clause to the transaction_date column.
**B.** Micro-partitions are fixed 500 MB storage units that guarantee even data distribution and are processed by dedicated compute clusters.
**C.** Micro-partitions maintain column metadata, including minimum and maximum values for transaction_date, allowing Snowflake to prune micro-partitions that cannot contain matching rows.

**D.** Micro-partitions automatically convert the transaction data into a row-based storage format that is faster for analytical queries.

**Answer: (SHOW ANSWER)**

Snowflake automatically divides table data into immutable, columnar micro-partitions. Users do not manually define these partitions.

Snowflake stores metadata for each micro-partition, including information such as the range of values for individual columns. When a query filters transaction_date to January 2025, the optimizer can compare the predicate with the minimum and maximum date values recorded for each micro-partition.

Micro-partitions whose date ranges cannot contain qualifying rows are excluded from the scan. This operation is called micro-partition pruning and can significantly reduce the amount of data read.

Micro-partitions normally contain approximately 50 MB to 500 MB of uncompressed data; they are not fixed at 500 MB.

Snowflake uses columnar-not row-based-storage for table data.

References:

Snowflake Documentation - Micro-partitions and data clustering: https://docs.snowflake.com/en/user-guide /tables-clustering-micropartitions

Snowflake Documentation - Micro-partition pruning: https://docs.snowflake.com/en/user-guide/tables- clustering-micropartitions#query-pruning

==

**NEW QUESTION: 31**

By default, which role can create resource monitors?

**A.** ACCOUNTADMIN

**B.** SECURITYADMIN

**C.** SYSADMIN

**D.** USERADMIN

**Answer: (SHOW ANSWER)**

The role that can by default create resource monitors in Snowflake is the ACCOUNTADMIN role. Resource monitors are a crucial feature in Snowflake that allows administrators to track and control the consumption of compute resources, ensuring that usage stays within specified limits. The creation and management of resource monitors involve defining thresholds for credits usage, setting up notifications, and specifying actions to be taken when certain thresholds are exceeded.

Given the significant impact that resource monitors can have on the operational aspects and billing of a Snowflake account, the capability to create and manage them is restricted to the ACCOUNTADMIN role. This role has the broadest set of privileges in Snowflake, including the ability to manage all aspects of the account, such as users, roles, warehouses, databases, and resource monitors, among others.

References:

Snowflake Documentation on Resource Monitors: Managing Resource Monitors

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

**NEW QUESTION: 32**
What is the SNOWFLAKE.ACCOUNT_USAGE view that contains information about which objects were read by queries within the last 365 days (1 year)?
**A.** VIEWS_HISTORY
**B.** OBJECT_HISTORY
**C.** ACCESS_HISTORY
**D.** LOGIN_HISTORY
**Answer: (SHOW ANSWER)**
The ACCESS_HISTORY view in the SNOWFLAKE.ACCOUNT_USAGE schema contains information about the access history of Snowflake objects, such as tables and views, within the last 365 days1.

**NEW QUESTION: 33**
Which kind of Snowflake table stores file-level metadata for each file in a stage?
**A.** Directory
**B.** External
**C.** Temporary
**D.** Transient
**Answer: (SHOW ANSWER)**
The kind of Snowflake table that stores file-level metadata for each file in a stage is a directory table.A directory table is an implicit object layered on a stage and stores file-level metadata about the data files in the stage3.

**NEW QUESTION: 34**
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: 35**
Which object provides a list of the files stored in a stage?
**A.** Semantic view

**B.** Directory table
**C.** Materialized view
**D.** Information schema
**Answer: (SHOW ANSWER)**
The correct answer is B. Directory table .
A directory table stores file-level metadata for files in a Snowflake stage. It can be used to see a list of files available in an internal or external stage.
Why B is correct:
Directory tables allow users to query staged file metadata, including file path, file size, last modified timestamp, and file URL information.
Example:
SELECT *
FROM DIRECTORY(@my_stage);
This returns metadata for files stored in the stage.
Why the other options are incorrect:
A). A semantic view is used to define business-friendly semantic relationships for analytics, not to list staged files.
C). A materialized view stores precomputed query results, not stage file metadata.
D). Information Schema contains metadata about Snowflake objects, but the Snowflake object specifically designed to list files in a stage is a directory table.
Official Snowflake documentation reference:
Snowflake documentation describes directory tables as objects that store a catalog of staged files in cloud storage.
Reference: Snowflake Documentation - Directory tables; Snowflake Documentation - Stages; SnowPro Core Study Guide - Data Loading and Unloading.
==

**NEW QUESTION: 36**
Using variables in Snowflake is denoted by using which SQL character?
**A.** @
**B.** &
**C.** $
**D.** #
**Answer: (SHOW ANSWER)**
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: 37**
What does SnowCD help Snowflake users to do?
**A.** Copy data into files.
**B.** Manage different databases and schemas.
**C.** Troubleshoot network connections to Snowflake.

**D.** Write SELECT queries to retrieve data from external tables.

**Answer: (SHOW ANSWER)**

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

**NEW QUESTION: 38**

What common query issues can be identified using the Query Profile? (Select TWO)

**A.** UNIONS without ALL

**B.** An incorrect filter condition in a WHERE clause

**C.** Insufficient role permissions

**D.** Exploding joins

**E.** Invalid SQL syntax

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 39**

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: 40**

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

**A.** account__usage . query history view

**B.** account__usage.query__history__archive View

**C.** information__schema . cruery_history view

**D.** information__schema - query history_by_ses s i on view

**Answer: (SHOW ANSWER)**

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

References:

[COF-C03] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Account Usage Schema1

**NEW QUESTION: 41**

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: 42**

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

**A.** Reader

**B.** Consumer

**C.** Vendor

**D.** Standard

**E.** Personalized

**Answer: (SHOW ANSWER)**

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

**NEW QUESTION: 43**

When unloading data, which combination of parameters should be used to differentiate between empty strings and NULL values? (Select TWO).

**A.** ESCAPE_UNENCLOSED_FIELD

**B.** REPLACE_INVALID_CHARACTERS

**C.** FIELD_OPTIONALLY_ENCLOSED_BY

**D.** EMPTY_FIELD_AS_NULL

**E.** SKIP_BLANK_LINES

**Answer: (SHOW ANSWER)**

When unloading data in Snowflake, it is essential to differentiate between empty strings and NULL values to preserve data integrity. The parameters FIELD_OPTIONALLY_ENCLOSED_BY and EMPTY_FIELD_AS_NULL are used together to address this:

FIELD_OPTIONALLY_ENCLOSED_BY: This parameter specifies the character used to enclose fields, which can differentiate between empty strings (as enclosed fields) and NULLs.

EMPTY_FIELD_AS_NULL: By setting this parameter, Snowflake interprets empty fields as NULL values when unloading data, ensuring accurate representation of NULLs versus empty strings.

These parameters are crucial when exporting data for systems that need explicit differentiation between NULL and empty string values.

**NEW QUESTION: 44**

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: 45**

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

**A.** The ability to scale up warehouses

**B.** The use of warehouse auto scaling

**C.** The ability to resize warehouses

**D.** Use of multi-clustered warehouses

**E.** The use of warehouse indexing

**Answer: (SHOW ANSWER)**

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

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

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

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

References:

Snowflake Documentation on Virtual Warehouses

SnowPro Core Certification Study Guide

**NEW QUESTION: 46**

What situation is likely to cause data spillage when a query is run?

**A.** When the query contains multiple filters and no data is returned.

**B.** When the number of queries exceeds the max_concurrency_level parameter setting.

**C.** When a virtual warehouse runs out of memory while executing the query.

**D.** When running queries that exceed the statement_timecut_in_seccnds parameter setting.

**Answer: C (LEAVE A REPLY)**

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

**NEW QUESTION: 47**
A clustering key was defined on a table, but It is no longer needed. How can the key be removed?
**A.** ALTER TABLE < TABLE NAME > PURGE CLUSTERING KEY
**B.** ALTER TABLE < TABLE NAME > DELETE CLUSTERING KEY
**C.** ALTER TABLE < TABLE NAME > DROP CLUSTERING KEY
**D.** ALTER TABLE < TABLE NAME > REMOVE CLUSTERING KEY
**Answer: (SHOW ANSWER)**
To remove a clustering key that was previously defined on a table in Snowflake, the correct SQL command isALTER TABLE < TABLE NAME > DROP CLUSTERING KEY. This command removes the existing clustering key from the table, after which Snowflake will no longer re-cluster data based on that key during maintenance operations or after data loading operations.
Reference: Snowflake Documentation on altering tables, specifically the section on modifying clustering keys: https://docs.snowflake.com/en/sql-reference/sql/alter-table.html#modifying-clustering-keys

**NEW QUESTION: 48**
Which functions are used to aggregate values into an array? Select TWO.
**A.** ARRAY_AGG
**B.** ARRAY_SLICE
**C.** ARRAY_APPEND
**D.** ARRAY_CONSTRUCT
**E.** ARRAY_TO_STRING
**Answer: (SHOW ANSWER)**
The clearly correct answer from the listed options is A. ARRAY_AGG .
Important correction:
The provided answer A, D is not fully correct if the question is asking about aggregate functions .
ARRAY_CONSTRUCT creates an array from explicitly supplied values, but it does not aggregate values across multiple rows.
Why A is correct:
ARRAY_AGG is an aggregate function. It aggregates input values from multiple rows into a single array.
Example:
SELECT ARRAY_AGG(product_id)
FROM orders;

This returns an array containing product_id values from multiple rows.
Why D is not an aggregate function:
ARRAY_CONSTRUCT creates an array from the values passed into the function.
Example:
SELECT ARRAY_CONSTRUCT( ' A ' , ' B ' , ' C ' );
This creates an array from listed expressions, but it does not aggregate values across rows.
Why the other options are incorrect:
B). ARRAY_SLICE returns a subset of an existing array.
C). ARRAY_APPEND appends a value to an existing array.
E). ARRAY_TO_STRING converts an array to a string.
Official Snowflake documentation reference:
Snowflake documentation describes ARRAY_AGG as an aggregate function that returns the input values, pivoted into an array. ARRAY_CONSTRUCT is documented as a scalar function that constructs an array from input expressions.
Reference: Snowflake Documentation - ARRAY_AGG; Snowflake Documentation - ARRAY_CONSTRUCT; SnowPro Core Study Guide - Working with Semi-Structured Data.
==

**NEW QUESTION: 49**
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: 50**
What Snowflake feature or service is used to prevent third parties from reading data while it is at rest or in transit?
Freecram.net
**A.** Zero-copy cloning
**B.** Network policies
**C.** Incremental data snapshots
**D.** End-to-end encryption
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 51**
Which of the following practices are recommended when creating a user in Snowflake? (Choose two.)
**A.** Set the user's access to expire within a specified timeframe.
**B.** Set the number of minutes to unlock to 15 minutes.

**C.** Force an immediate password change.
**D.** Set a default role for the user.
**E.** Configure the user to be initially disabled.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 52**
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: 53**
Which type of workload is recommended for Snowpark-optimized virtual warehouses?
**A.** Workloads with ad hoc analytics
**B.** Workloads that have large memory requirements
**C.** Workloads with unpredictable data volumes for each query
**D.** Workloads that are queried with small table scans and selective filters
**Answer: (SHOW ANSWER)**
Snowpark-optimized virtual warehouses in Snowflake are designed to efficiently handle workloads with large memory requirements. Snowpark is a developer framework that allows users to write code in languages like Scala, Java, and Python to process data in Snowflake. Given the nature of these programming languages and the types of data processing tasks they are typically used for, having a virtual warehouse that can efficiently manage large memory-intensive operations is crucial.
Understanding Snowpark-Optimized Virtual Warehouses:
Snowpark allows developers to build complex data pipelines and applications within Snowflake using familiar programming languages.
These virtual warehouses are optimized to handle the execution of Snowpark workloads, which often involve large datasets and memory-intensive operations.
Large Memory Requirements:
Workloads with large memory requirements include data transformations, machine learning model training, and advanced analytics.
These operations often need to process significant amounts of data in memory to perform efficiently.
Snowpark-optimized virtual warehouses are configured to provide the necessary memory resources to support these tasks, ensuring optimal performance and scalability.
Other Considerations:
While Snowpark can handle other types of workloads, its optimization for large memory tasks makes it particularly suitable for scenarios where data processing needs to be done in-memory.

Snowflake's ability to scale compute resources dynamically also plays a role in efficiently managing large memory workloads, ensuring that performance is maintained even as data volumes grow.
Snowflake Documentation: Introduction to Snowpark
Snowflake Documentation: Virtual Warehouses

**NEW QUESTION: 54**
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: 55**
A user received a set of PDF files that are available on an external stage. Text, data, and layout elements need to be extracted from the files. Which Snowflake AI function should be used?
**A.** TRANSCRIBE
**B.** COMPLETE
**C.** PARSE_DOCUMENT
**D.** EMBED
**Answer: (SHOW ANSWER)**
PARSE_DOCUMENT is used to extract information from documents, including text and layout elements. It is part of Snowflake Cortex document processing functionality.
Why the other options are incorrect:
TRANSCRIBE is for audio transcription, COMPLETE is for large language model completion, and EMBED generates vector embeddings.
Reference: Snowflake Documentation - Cortex PARSE_DOCUMENT; Document AI; Snowflake Cortex.
==

**NEW QUESTION: 56**
A user has a standard multi-cluster warehouse auto-scaling policy in place.
Which condition will trigger a cluster to shut-down?
**A.** When after 2-3 consecutive checks the system determines that the load on the most-loaded cluster could be redistributed.
**B.** When after 5-6 consecutive checks the system determines that the load on the most-loaded cluster could be redistributed.

**C.** When after 5-6 consecutive checks the system determines that the load on the least-loaded cluster could be redistributed.
**D.** When after 2-3 consecutive checks the system determines that the load on the least-loaded cluster could be redistributed.
**Answer: D (LEAVE A REPLY)**
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: 57**
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: 58**
Which task privilege does a Snowflake role need in order to suspend or resume a task?
**A.** USAGE
**B.** OPERATE
**C.** MONITOR
**D.** OWNERSHIP
**Answer: (SHOW ANSWER)**
In Snowflake, the OPERATE privilege is required for a role to suspend or resume a task.This privilege allows the role to perform operational tasks such as starting and stopping tasks, which includes suspending and resuming them6

**NEW QUESTION: 59**
What is the minimum Snowflake edition required to create a materialized view?
**A.** Standard Edition
**B.** Enterprise Edition
**C.** Business Critical Edition
**D.** Virtual Private Snowflake Edition

**Answer: (SHOW ANSWER)**

Materialized views in Snowflake are a feature that allows for the pre-computation and storage of query results for faster query performance. This feature is available starting from the Enterprise Edition of Snowflake. It is not available in the Standard Edition, and while it is also available in higher editions like Business Critical and Virtual Private Snowflake, the Enterprise Edition is the minimum requirement.

References:

Snowflake Documentation on CREATE MATERIALIZED VIEW1.

Snowflake Documentation on Working withMaterialized Views

https://docs.snowflake.com/en/sql-reference/sql/create-materialized-view.html#:~:text=Materialized%20views%20require%20Enterprise%20Edition,upgrading%2C%20please%20contact%20Snowflake%20Support.

**NEW QUESTION: 60**

By default, how long is the standard retention period for Time Travel across all Snowflake accounts?

**A.** 0 days

**B.** 1 day

**C.** 7 days

**D.** 14 days

**Answer: (SHOW ANSWER)**

By default, the standard retention period for Time Travel in Snowflake is 1 day across all Snowflake accounts. Time Travel enables users to access historical data within this retention window, allowing for point- in-time data analysis and recovery. This feature is a significant aspect of Snowflake ' s data management capabilities, offering flexibility in handling data changes and accidental deletions.

References:

Snowflake Documentation: Using Time Travel

**NEW QUESTION: 61**

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

**A.** True

**B.** False

**Answer: B (LEAVE A REPLY)**

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

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

**Valid COF-C03 Dumps** shared by ExamDiscuss.com for Helping Passing COF-C03 Exam! ExamDiscuss.com now offer the **newest COF-C03 exam dumps**, the ExamDiscuss.com COF-C03 exam **questions have been updated**

and **answers have been corrected** get the **newest** ExamDiscuss.com COF-C03 dumps with Test Engine here: https://www.examdiscuss.com/Snowflake/exam/COF-C03/premium/ (**1187** Q&As Dumps, **35%OFF** **Special Discount Code:** **freecram**)

**NEW QUESTION: 62**
Which view should be checked to ensure that data has loaded successfully before removing staged files?
**A.** STAGE_STORAGE_USAGE_HISTORY
**B.** LOAD_HISTORY
**C.** COPY_HISTORY
**D.** DATA TRANSFER HISTORY
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 63**
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 isOPERATE.
TheOPERATEprivilege 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: 64**
Which pages are included in the Activity area of Snowsight? (Select TWO).
**A.** Contacts
**B.** Sharing settings
**C.** Copy History
**D.** Query History
**E.** Automatic Clustering History
**Answer: (SHOW ANSWER)**
The Activity area of Snowsight includes the Query History page, which allows users to monitor and view details about queries executed in their account, including performance data1.It also includes the Automatic Clustering History, which provides insights into the automatic clustering operations performed on tables2.

**NEW QUESTION: 65**
Which feature is only available in the Enterprise or higher editions of Snowflake?
**A.** Column-level security

**B.** SOC 2 type II certification
**C.** Multi-factor Authentication (MFA)
**D.** Object-level access control
**Answer: A (LEAVE A REPLY)**
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: 66**
What is the FASTEST way to load a specific, staged, discrete set of files using the COPY INTO < table > command?
**A.** Use pattern matching using the PATTERN parameter.
**B.** Select which files to load by defining the path prefix on an external stage.
**C.** Use wildcard characters in the file path definition.
**D.** List the file names using the FILES parameter.
**Answer: (SHOW ANSWER)**
The fastest way to load a specific, discrete set of staged files is to explicitly list the file names using the FILES parameter in the COPY INTO < table > command.
Snowflake recommends using the FILES parameter when you already know the exact files that need to be loaded. This avoids the overhead of scanning and matching files using a regular expression pattern.
Why D is correct:
The FILES parameter lets you specify an exact list of staged files to load. Since Snowflake does not need to evaluate a pattern against all available staged files, this is the fastest method for loading a known, specific set of files.
Why the other options are incorrect:
A). PATTERN uses regular expression matching, which can be useful, but it is generally slower than listing specific files with FILES.
B). A path prefix can narrow the file location, but it does not explicitly identify a discrete set of files.
C). Wildcard characters are not the recommended fastest method for loading a known set of staged files in Snowflake.
Official Snowflake documentation reference:
Snowflake COPY INTO < table > documentation states that the FILES parameter is used to specify a list of one or more files to load, and Snowflake recommends using FILES instead of PATTERN for better performance when loading a specific list of files.
Reference: Snowflake Documentation - COPY INTO < table > command, FILES and PATTERN parameters.

**NEW QUESTION: 67**
Which data types are supported by Snowflake when using semi-structured data? (Choose two.)
**A.** VARIANT
**B.** VARRAY

**C.** STRUCT
**D.** ARRAY
**E.** QUEUE
**Answer: A,D (LEAVE A REPLY)**
Snowflake supports the VARIANT and ARRAY data types for semi-structured data. VARIANT can store values of any other type, including OBJECT and ARRAY, making it suitable for semi-structured data formats like JSON. ARRAY is used to store an ordered list of elements

**NEW QUESTION: 68**
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: 69**
What is cached during a query on a virtual warehouse?
**A.** All columns in a micro-partition
**B.** Any columns accessed during the query
**C.** The columns in the result set of the query
**D.** All rows accessed during the query
**Answer: (SHOW ANSWER)**
During a query on a virtual warehouse, the columns in the result set of the query are cached. This allows for faster retrieval of data if the same or a similar query is run again, as the system can retrieve the data from the cache rather than reprocessing the entire query. References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 70**
What is used to denote a pre-computed data set derived from a SELECT query specification and stored for later use?
Freecram.net
**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: 71**
Which methods for a second factor of authentication does Snowflake support? (Choose two.)
**A.** Email codes
**B.** Authenticator applications
**C.** Social logins
**D.** Passkeys
**E.** Security questions
**Answer: (SHOW ANSWER)**
Snowflake supports multiple methods for a second authentication factor, including passkeys, authenticator applications that generate time-based one-time passcodes, and Duo. Among the listed options, authenticator applications and passkeys are both supported. A passkey uses public-key cryptography and can be stored on a device, hardware security key, or password manager. An authenticator application produces a time-based one-time passcode that the user enters after the password. Email codes, social logins, and security questions are not the supported MFA methods represented in the choices. Administrators can use authentication policies to control which MFA methods users may enroll and use. Because Duo is supported but is not listed as an option, the two correct answers from the available choices are B and D.

**NEW QUESTION: 72**
Two users share a virtual warehouse named wh dev 01. When one of the users loads data, the other one experiences performance issues while querying data.
How does Snowflake recommend resolving this issue?
**A.** Scale up the existing warehouse.
**B.** Create separate warehouses for each user.
**C.** Create separate warehouses for each workload.
**D.** Stop loading and querying data at the same time.
**Answer: (SHOW ANSWER)**
Snowflake recommends creating separate warehouses for each workload to resolve performance issues caused by shared virtual warehouses. This ensures that the resources are not being overutilized by one user's activities, thereby affecting the performance of another user's activities4.

**NEW QUESTION: 73**
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: 74**
A user needs to implement a Slowly Changing Dimension Type 2 to track historical changes.
Which action will accomplish this?
**A.** Enable CHANGE_TRACKING on the dimension table and use a scheduled task to run a MERGE that reads only changed rows.
**B.** Use Time Travel with periodic clones to snapshot the dimension and reconstruct its history using views.
**C.** Use Snowpipe for automatic ingestion and materialized views to propagate changes into the dimension table for versioning.
**D.** Use a stream on the source or staging table with a task that runs a MERGE to expire current rows and insert new versioned rows with effective dates.
**Answer: (SHOW ANSWER)**
A Type 2 Slowly Changing Dimension preserves history by expiring the currently active dimension record and inserting a new version whenever tracked attributes change. Snowflake recommends streams and tasks for this requirement because a stream records DML changes from the source or staging table, and a scheduled or triggered task can process those changes with a MERGE statement.
The MERGE can update the old row's end date or current-row flag and insert a replacement row with new effective dates. Dynamic tables represent a maintained current result and are not the preferred solution for Type 2 historical tracking. Time Travel and clones are recovery and historical-query features, while Snowpipe and materialized views do not implement the required record-versioning workflow.

**NEW QUESTION: 75**
Which of the following can be used when unloading data from Snowflake? (Choose two.)
**A.** When unloading semi-structured data, it is recommended that the STRIP_OUTER_ARRAY option be used.
**B.** Use the ENCODING file format option to change the encoding from the default UTF-8.
**C.** The OBJECT_CONSTRUCT function can be used to convert relational data to semi-structured data.
**D.** By using the SINGLE = TRUE parameter, a single file up to 5 GB in size can be exported to the storage layer.
**E.** Use the PARSE_JSON function to ensure structured data will be unloaded into the VARIANT data type.
**Answer: (SHOW ANSWER)**
The OBJECT_CONSTRUCT function is used in Snowflake to create a JSON object from relational data, which is useful when unloading semi-structured data. The SINGLE = TRUE parameter is used when unloading data to ensure that the data is exported as a single file, which can be up to 5 GB in size. References:
[COF-C03] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 76**
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: D (LEAVE A REPLY)**
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

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

**NEW QUESTION: 77**
How can a user access information about a query execution plan without consuming virtual warehouse compute resources?
**A.** Use the Snowsight dashboard.
**B.** Review the data in the Account_Usage view.
**C.** Review the Query Profile metrics.
**D.** Use the EXPLAIN function.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 78**
How can a Snowsight user change a Standard virtual warehouse to a Snowpark-optimized virtual warehouse?
**A.** Use the ALTER WAREHOUSE command on an active Standard virtual warehouse
**B.** Use the alter warehouse command on an active Snowpark-optimized warehouse.
**C.** Use the ALTER warehouse command on a suspended Standard virtual warehouse.
**D.** Use the ALTER WAREHOUSE command on a suspended Snowpark-optimized warehouse.
**Answer: (SHOW ANSWER)**
To change a Standard virtual warehouse to a Snowpark-optimized virtual warehouse, the warehouse must be in a suspended state. This ensures that no operations are disrupted during the modification process.

Suspend the Warehouse: Ensure that the warehouse is suspended.
ALTER WAREHOUSE my_warehouse SUSPEND;
Alter the Warehouse: Modify the warehouse to be Snowpark-optimized.
ALTER WAREHOUSE my_warehouse SET WAREHOUSE_TYPE = 'SNOWPARK-OPTIMIZED'; Resume the Warehouse: Resume the warehouse to make it operational.
ALTER WAREHOUSE my_warehouse RESUME;
References:
Snowflake Documentation: Creating Snowpark-Optimized Warehouses
Snowflake Documentation: ALTER WAREHOUSE

**NEW QUESTION: 79**
How should a Snowflake use' configure a virtual warehouse to be in Maximized mode"
**A.** Set the WAREHOUSES_SIZE to 6XL.
**B.** Set the STATEMENT_TIMEOUT_1M_SECOMES to 0.
**C.** Set the MAX_CONCURRENCY_LEVEL to a value of 12 or large.
**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: 80**
User1, who has the SYSADMIN role, executed a query on Snowsight. User2, who is in the same Snowflake account, wants to view the result set of the query executed by User1 using the Snowsight query history.
What will happen if User2 tries to access the query history?
Freecram.net
**A.** If User2 has the sysadmin role they will be able to see the results.
**B.** If User2 has the securityadmin role they will be able to see the results.
**C.** If User2 has the ACCOUNTADMIN role they will be able to see the results.
**D.** User2 will be unable to view the result set of the query executed by User1.
**Answer: C (LEAVE A REPLY)**
In Snowflake, the query history and the results of queries executed by a user are accessible based on the roles and permissions. If User1 executed a query with the SYSADMIN role, User2 would be able to view the result set of that query executed by User1 only if User2 has the ACCOUNTADMIN role. The ACCOUNTADMIN role has the broadest set

of privileges, including the ability to access all aspects of the account ' s operation, data, and query history, thus enabling User2 to view the results of queries executed by other users.
References:
Snowflake Documentation: Understanding Snowflake Roles

**NEW QUESTION: 81**
A Medium-sized standard virtual warehouse is being used to load data continuously. The data is also consumed through reports.
Which step will optimize costs?
**A.** Create separate warehouses for each workload.
**B.** Enable the Query Acceleration Service.
**C.** Change to a multi-cluster warehouse.
**D.** Resize the warehouse to Small.
**Answer: A (LEAVE A REPLY)**
Using separate virtual warehouses isolates the continuous loading workload from reporting queries.
Each warehouse can be sized, monitored, suspended, and resumed independently, so reporting activity does not compete with ingestion operations for the same compute resources. This also allows the loading warehouse and reporting warehouse to be right-sized for their different workload characteristics and cost profiles. Multi-cluster warehouses primarily address query concurrency and are not intended to improve data-loading performance. Query Acceleration Service helps eligible analytical queries but is not a general loading-cost optimization. Simply resizing the shared warehouse to Small could lengthen both loading and report execution and might increase total runtime. Therefore, separating the workloads is the most appropriate cost and performance optimization among the available choices.

**NEW QUESTION: 82**
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: 83**
What are best practice recommendations for using the ACCOUNTADMIN system-defined role in Snowflake? (Choose two.)
**A.** Ensure all ACCOUNTADMIN roles use Multi-factor Authentication (MFA).
**B.** All users granted ACCOUNTADMIN role must be owned by the ACCOUNTADMIN role.
**C.** The ACCOUNTADMIN role must be granted to only one user.
**D.** Assign the ACCOUNTADMIN role to at least two users, but as few as possible.
**E.** All users granted ACCOUNTADMIN role must also be granted SECURITYADMIN role.
**Answer: (SHOW ANSWER)**
Best practices for using the ACCOUNTADMIN role include ensuring that all users with this role use Multi- factor Authentication (MFA) for added security. Additionally, it is recommendedto assign the ACCOUNTADMIN role to at least two users to avoid delays in case of password recovery issues, but to as few users as possible to maintain strict control over account-level operations4.

**NEW QUESTION: 84**
Which commands will indicate if Multi-Factor Authentication, or MFA, has been enabled for a given user?
Select TWO.
**A.** SHOW USERS
**B.** DESCRIBE USER
**C.** SHOW PARAMETERS
**D.** SHOW USER FUNCTIONS
**E.** SHOW CONNECTIONS
**Answer: (SHOW ANSWER)**
The correct answers are A. SHOW USERS and B. DESCRIBE USER .
Snowflake user metadata includes information related to MFA enrollment and configuration. To inspect user- level properties, administrators can use SHOW USERS or DESCRIBE USER.
Why A is correct:
SHOW USERS returns information about users in the account. Its output can include user properties that indicate MFA status or related authentication details.
Why B is correct:

DESCRIBE USER returns detailed properties for a specific user. This can be used to inspect whether MFA has been enabled for that user.
Example:
DESCRIBE USER my_user;
Why the other options are incorrect:
C). SHOW PARAMETERS displays parameters, not user MFA enrollment status.
D). SHOW USER FUNCTIONS lists user-defined functions, not Snowflake users.
E). SHOW CONNECTIONS displays connection objects and does not indicate MFA status for a specific user.
Official Snowflake documentation reference:
Snowflake documentation describes SHOW USERS and DESCRIBE USER as commands for viewing user metadata and user properties.
Reference: Snowflake Documentation - SHOW USERS; DESCRIBE USER; Multi-Factor Authentication; SnowPro Core Study Guide - Security and Access Control.
==

**NEW QUESTION: 85**
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: B (LEAVE A REPLY)**
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: 86**
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: 87**

How does the search optimization service help Snowflake users improve query performance?
**A.** It scans the micro-partitions based on the joins used in the queries and scans only join columns.
**B.** Il maintains a persistent data structure that keeps track of the values of the table ' s columns m each of its micro-partitions.
**C.** It scans the local disk cache to avoid scans on the tables used in the Query.
**D.** It keeps track of running queries and their results and saves those extra scans on the table.
**Answer: (SHOW ANSWER)**
The search optimization service in Snowflake enhances query performance by maintaining a persistent data structure. This structure indexes the values of table columns across micro-partitions, allowing Snowflake to quickly identify which micro-partitions contain relevant data for a query. By efficiently narrowing down the search space, this service reduces the amount of data scanned during query execution, leading to faster response times and more efficient use of resources.
References: Snowflake Documentation on Search Optimization Service

**NEW QUESTION: 88**
Which SQL statement will require a virtual warehouse to run?
**A.** SELECT COUNT{*) FROM TBL_EMPLOYEE;
**B.** ALTER TABLE TBL_EMPLOYEE ADD COLUMN EMP_REGI0N VARCHAR(20);
**C.** INSERT INTO TBL_EMPLOYEE(EMP_ID, EMP_NAME, EMP_SALARY, DEPT) VALUES(1, ' Adam*,20000,* Finance ' );
**D.** CREATE OR REPLACE TABLE TBL_EMPLOYEE (
EMP_ID NUMBER,
EMP_NAME VARCHAR(30),
EMP_SALARY NUMBER, DEPT VARCHAR{20)
);
**Answer: (SHOW ANSWER)**
A virtual warehouse in Snowflake is used to perform data processing tasks that require computational resources, such as queries that modify data or perform significant computation. Of the options provided:
C). INSERT INTO TBL_EMPLOYEE(EMP_ID, EMP_NAME, EMP_SALARY, DEPT) VALUES(1, ' Adam
' ,20000, ' Finance ' );This SQL statement performs a data modification operation (DML) by inserting a new record into theTBL_EMPLOYEEtable, which requires computational resources provided by a virtual warehouse to execute.
References:
Snowflake Documentation: Understanding Virtual Warehouses

**NEW QUESTION: 89**
Query compilation occurs in which architecture layer of the Snowflake Cloud Data Platform?
**A.** Compute layer
**B.** Storage layer
**C.** Cloud infrastructure layer
**D.** Cloud services layer
**Answer: (SHOW ANSWER)**

Query compilation in Snowflake occurs in the Cloud Services layer. This layer is responsible for coordinating and managing all aspects of the Snowflake service, including authentication, infrastructure management, metadata management, query parsing and optimization, and security. By handling these tasks, the Cloud Services layer enables the Compute layer to focus on executing queries, while the Storage layer is dedicated to persistently storing data.
References:
[COF-C03] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Snowflake Architecture1

**NEW QUESTION: 90**
Which queries will require a running virtual warehouse? (Select TWO).
**A.** SELECT CURRENT_DATABASE();
**B.** DROP TABLE < table > ;
**C.** UPDATE < table > SET;
**D.** DELETE FROM < table > ;
**E.** SELECT COUNT(*) FROM < table > ;
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 91**
Users who do not have a Snowflake account need access to a staged file. Which command will generate a URL that will provide access to those users?
**A.** BUILD_STAGE_FILE_URL(@ < stage_name > , ' < relative_file_path > ' );
**B.** GET_PRESIGNED_URL(@ < stage_name > , ' < relative_file_path > ' , [ < expiration_time > ]);
**C.** GET_RELATIVE_PATH(@ < stage_name > , ' < absolute_file_path > ' );
**D.** BUILD_SCOPED_FILE_URL(@ < stage_name > , ' < relative_file_path > ' );
**Answer: (SHOW ANSWER)**
The correct answer is B. GET_PRESIGNED_URL(@ < stage_name > , ' < relative_file_path > ' , [ < expiration_time > ]); .
A pre-signed URL provides temporary access to a staged file. This is useful when users or applications outside Snowflake need access to a file for a limited period of time.
Why B is correct:
GET_PRESIGNED_URL generates a pre-signed URL for a staged file. The URL can be used to access the file without the user needing a Snowflake account, as long as the URL has not expired.
Example:
SELECT GET_PRESIGNED_URL(@my_stage, ' folder/file.csv ' , 3600);
This generates a URL that is valid for a limited time, such as 3600 seconds.
Why the other options are incorrect:
A). BUILD_STAGE_FILE_URL creates a Snowflake file URL, but it does not provide the same temporary external access behavior as a pre-signed URL.
C). GET_RELATIVE_PATH returns the relative path of a staged file, not an access URL.
D). BUILD_SCOPED_FILE_URL creates a scoped URL that is useful within Snowflake-controlled contexts, but for users without Snowflake accounts who need direct temporary access, GET_PRESIGNED_URL is the best answer.

Official Snowflake documentation reference:
Snowflake documentation describes GET_PRESIGNED_URL as a function that returns a pre-signed URL for a staged file, with an optional expiration time.
Reference: Snowflake Documentation - GET_PRESIGNED_URL; Snowflake Documentation - File functions; SnowPro Core Study Guide - Data Loading and Unloading.
==

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

**NEW QUESTION: 92**
What COPY INTO SQL command should be used to unload data into multiple files?
**A.** SINGLE=TRUE
**B.** MULTIPLE=TRUE
**C.** MULTIPLE=FALSE
**D.** SINGLE=FALSE
**Answer: (SHOW ANSWER)**
The COPY INTO SQL command with the option SINGLE=FALSE is used to unload data into multiple files.
This option allows the data to be split into multiple files during the unload process. References: SnowPro Core Certification COPY INTO SQLcommand unload multiple files

**NEW QUESTION: 93**
A Snowflake user executed a query and received the results. Another user executed the same query 4 hours later. The data had not changed.
What will occur?
**A.** No virtual warehouse will be used, data will be read from the result cache.
**B.** No virtual warehouse will be used, data will be read from the local disk cache.
**C.** The default virtual warehouse will be used to read all data.
**D.** The virtual warehouse that is defined at the session level will be used to read all data.
**Answer: (SHOW ANSWER)**
Snowflake maintains a result cache that stores the results of every query for 24 hours.If the same query is executed again within this time frame and the data has not changed, Snowflake will retrieve the data from the result cache instead of using a virtual warehouse to recompute the results2.

**NEW QUESTION: 94**

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: (SHOW ANSWER)**

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: 95**

Which built-in function should be used to determine whether the current role has inherited privileges from another specified role?

**A.** INVOKER_ROLE

**B.** CURRENT_ROLE

**C.** IS_ROLE_IN_SESSION

**D.** CURRENT_ROLE_TYPE

**Answer: (SHOW ANSWER)**

IS_ROLE_IN_SESSION determines whether the user's active primary or secondary role hierarchy inherits the privileges of a specified account role. It returns TRUE when the specified role is active or lies below an active role in the hierarchy, and FALSE otherwise. This makes it useful in row access policies, masking policies, and other conditional security logic where inherited privileges must be considered. CURRENT_ROLE returns only the name of the current primary role and does not evaluate inheritance. INVOKER_ROLE identifies the role that invoked an object in the relevant execution

context, while CURRENT_ROLE_TYPE reports the type of the current role. The source option's incorrect dash has been corrected to the valid Snowflake function name, IS_ROLE_IN_SESSION.

**NEW QUESTION: 96**
Revoking which privilege will prevent one role from giving privileges to other roles?
**A.** USAGE
**B.** MODIFY
**C.** REFERENCES
**D.** MANAGE GRANTS
**Answer: (SHOW ANSWER)**
The correct answer is D. MANAGE GRANTS .
The MANAGE GRANTS privilege allows a role to grant and revoke privileges on objects. If this privilege is revoked, the role loses the broad ability to manage privileges for other roles.
Why D is correct:
MANAGE GRANTS is a powerful global privilege. It allows a role to grant and revoke privileges, even on objects the role does not own, depending on Snowflake's access control rules.
Why the other options are incorrect:
A). USAGE allows access to an object or namespace, such as a database, schema, warehouse, or integration. It does not allow broad grant management.
B). MODIFY allows changes to certain object properties or data, depending on the object type, but it does not broadly allow granting privileges to roles.
C). REFERENCES is related to referencing objects, such as in constraints or policies. It does not control grant administration.
Official Snowflake documentation reference:
Snowflake documentation describes MANAGE GRANTS as the global privilege that enables granting and revoking privileges on objects.
Reference: Snowflake Documentation - Access control privileges; Snowflake Documentation - MANAGE GRANTS; SnowPro Core Study Guide - Security and Access Control.
==

**NEW QUESTION: 97**
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: (SHOW ANSWER)**

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: 98**
Which command will unload data from a table into an external stage?
**A.** PUT
**B.** INSERT
**C.** COPY INTO < location >
**D.** GET
**Answer: (SHOW ANSWER)**
In Snowflake, theCOPY INTO < location > command is used to unload (export) data from a Snowflake table to an external stage, such as an S3 bucket, Azure Blob Storage, or Google Cloud Storage. This command allows users to specify the format, file size, and other options for the data being unloaded, making it a flexible solution for exporting data from Snowflake to external storage solutions for further use or analysis.
References: Snowflake Documentation on Data Unloading

**NEW QUESTION: 99**
A user has 10 files in a stage containing new customer data. The ingest operation completes with no errors, using the following command:
COPY INTO my__table FROM @my__stage;
The next day the user adds 10 files to the stage so that now the stage contains a mixture of new customer data and updates to the previous data. The user did not remove the 10 original files.
If the user runs the same copy into command what will happen?
**A.** All data from all of the files on the stage will be appended to the table
**B.** Only data about new customers from the new files will be appended to the table
**C.** The operation will fail with the error uncertain files in stage.
**D.** All data from only the newly-added files will be appended to the table.
**Answer: (SHOW ANSWER)**
When the COPY INTO command is executed in Snowflake, it processes all files present in the specified stage that have not been ingested before or marked as already loaded. Since the user did not remove the original 10 files after the first load, running the same COPY INTO command again will result in all 20 files being processed. This means that the data from the original 10 files will be appended to the table again, along with the data from the new 10 files, potentially leading to duplicate records for the original data set.
References:
Snowflake Documentation on Data Loading
SnowPro Core Certification Study Guide

**NEW QUESTION: 100**
Which action can be performed by the SYSADMIN role without requiring elevated privileges from the ACCOUNTADMIN or SECURITYADMIN role?

Freecram.net

**A.** Granting privileges
**B.** Creating a new user
**C.** Managing virtual warehouses
**D.** Viewing billing and usage information
**Answer: (SHOW ANSWER)**
The SYSADMIN system role is intended to manage Snowflake compute and database objects. By default, it can create virtual warehouses, databases, schemas, and other database objects, and it can manage objects it owns or has been granted privileges on. Creating users belongs to the USERADMIN role hierarchy, while broad privilege administration is handled by SECURITYADMIN through MANAGE GRANTS. Billing and account-usage administration normally require ACCOUNTADMIN or explicitly granted application or database roles and privileges. Therefore, managing virtual warehouses is the action that SYSADMIN can perform without first elevating to ACCOUNTADMIN or SECURITYADMIN. This separation supports least privilege by keeping object and compute administration separate from identity management, security-grant administration, and billing oversight.

**NEW QUESTION: 101**
Which file function provides a URL with access to a file on a stage without the need for authentication and authorization?
**A.** GET_RELATIVE_PATH
**B.** GET_PRESIGNED_URL
**C.** BUILD_STAGE_FILE_URL
**D.** BUILD_SCOPED_FILE_URL
**Answer: (SHOW ANSWER)**
TheGET_PRESIGNED_URLfile function in Snowflake provides a URL with access to a file on a stage without requiring authentication and authorization. This is particularly useful for sharing data files stored in Snowflake stages with external parties securely and conveniently. The presigned URL generated by this function gives temporary access to the file, which expires after a specified duration.
Example usage ofGET_PRESIGNED_URL:
SELECTGET_PRESIGNED_URL( ' < stage_name > ' , ' < file_path > ' );
This function generates a URL that can be used to directly access a file in the stage, making it easier to share data without compromising security.
Reference: Snowflake Documentation on File and Table Functions, specifically the section on GET_PRESIGNED_URL: https://docs.snowflake.com/en/sql-reference/functions/get_presigned_url.html Top of Form

**NEW QUESTION: 102**
What occurs when a pipe is recreated using the CREATE OR REPLACE PIPE command?
**A.** The Pipe load history is reset to empty.
**B.** The REFRESH command is executed.
**C.** The stage will be purged.
**D.** The destination table is truncated.
**Answer: (SHOW ANSWER)**

When a pipe is recreated using the CREATE OR REPLACE PIPE command, the load history of the pipe is reset. This means that Snowpipe will consider all files in the stage as new and will attempt to load them, evenif they were loaded previously by the old pipe2.

**NEW QUESTION: 103**
A user unloaded a Snowflake table called mytable to an internal stage called mystage.
Which command can be used to view the list of files that has been uploaded to the staged?
**A.** list @mytable;
**B.** list @%raytable;
**C.** list @ %m.ystage;
**D.** list @mystage;
**Answer: (SHOW ANSWER)**
The command list @mystage; is used to view the list of files that have been uploaded to an internal stage in Snowflake. The list command displays the metadata for all files in the specified stage, which in this case is mystage. This command is particularly useful for verifying that files have been successfully unloaded from a Snowflake table to the stage and for managing the files within the stage.
References:
Snowflake Documentation on Stages
SnowPro Core Certification Study Guide

**NEW QUESTION: 104**
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 caseof 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

**NEW QUESTION: 105**
When using the COPY INTO < location > command, which parameter is used to unload table rows into separate files?
**A.** DETAILED_OUTPUT
**B.** PARTITION BY
**C.** FILE_FORMAT
**D.** HEADER

**Answer: B (LEAVE A REPLY)**
The correct answer is B. PARTITION BY .
The PARTITION BY option in COPY INTO < location > partitions unloaded data into separate files based on the expression specified. This is commonly used to organize unloaded files by values such as date, region, or another column.
Why B is correct:
PARTITION BY controls how unloaded rows are grouped into different output paths/files.
Example:
COPY INTO @my_stage/sales/
FROM sales
PARTITION BY (TO_DATE(sale_date))
FILE_FORMAT = (TYPE = CSV);
This unloads rows into separate output paths based on the sale_date expression.
Why the other options are incorrect:
A). DETAILED_OUTPUT controls the amount of unload result information returned by the command.
C). FILE_FORMAT controls the format of the unloaded files, such as CSV, JSON, or Parquet.
D). HEADER controls whether column headings are included in unloaded output files.
Official Snowflake documentation reference:
Snowflake documentation for COPY INTO < location > describes PARTITION BY as the option used to partition unloaded rows into separate files based on an expression.
Reference: Snowflake Documentation - COPY INTO < location > ; Snowflake Documentation - Unloading data with partitioning; SnowPro Core Study Guide - Data Loading and Unloading.
==

**NEW QUESTION: 106**
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

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

**NEW QUESTION: 107**
What objects are required to unload data into cloud data storage? (Select TWO).
**A.** A directory table
**B.** An external table
**C.** A FILE FORMAT definition
**D.** A STORAGE_INTEGRATION definition
**E.** An external stage
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 108**
Which command is used in Snowflake to manually refresh a directory table?
**A.** TRUNCATE STAGE
**B.** ALTER STAGE
**C.** ALTER TABLE
**D.** ALTER VIEW
**Answer: (SHOW ANSWER)**
The correct answer is B. ALTER STAGE .
Directory tables are associated with stages. To manually refresh the metadata in a directory table, the stage is refreshed using the ALTER STAGE ... REFRESH command.
Why B is correct:
A directory table stores file-level metadata for files in a stage. Since the directory table is tied to the stage, the manual refresh operation is performed through the stage.
Example:
ALTER STAGE my_stage REFRESH;
This refreshes the directory table metadata for the stage.
Why the other options are incorrect:
A). TRUNCATE STAGE is not the command used to refresh directory table metadata.
C). ALTER TABLE is used to modify table properties, not refresh a directory table associated with a stage.
D). ALTER VIEW is used to modify views, not directory tables.
Official Snowflake documentation reference:
Snowflake documentation explains that directory table metadata can be refreshed manually using ALTER STAGE ... REFRESH.

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

**NEW QUESTION: 109**
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.** COPY INTO my_table FROM @many_stage PATTERN = '.\.csv';
**B.** COPY INTO my_table FROM @many_stage/app_files FILE_FORMAT = (TYPE = CSV);
**C.** COPY INTO my_table FROM @many_stage;
**D.** COPY INTO my_table FROM @many_stage/app_files/.csv;
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 110**
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: 111**

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 join operations

**B.** A column that appears most frequently in order by operations

**C.** A column that appears most frequently in where operations

**D.** A column that appears most frequently in aggregate operations

**E.** A column that appears most frequently in group by operations

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 112**

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

**A.** When creating a table

**B.** When loading data into a table

**C.** When unloading data from a table

**D.** When executing a show command

**E.** When executing a list command

**Answer: (SHOW ANSWER)**

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

**NEW QUESTION: 113**

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

**A.** Enterprise

**B.** Business Critical

**C.** Standard

**D.** Virtual Private Snowflake

**Answer: (SHOW ANSWER)**

Periodic rekeying of micro-partitions is a feature that requires the Enterprise Edition of Snowflake or higher. This feature is part of Snowflake's comprehensive approach to encryption key management, ensuring data security through regular rekeying1. References: [COF-C03] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 114**

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: 115**
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: 116**
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: 117**
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: A,C (LEAVE A REPLY)**
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: 118**
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: 119**
Awarding a user which privileges on all virtual warehouses is equivalent to granting the user the global MANAGE WAREHOUSES privilege?
**A.** MODIFY, MONITOR and OPERATE privileges
**B.** ownership and usage privileges
**C.** APPLYBUDGET and audit privileges
**D.** MANAGE LISTING ADTOTOLFillment and resolve all privileges
**Answer: A (LEAVE A REPLY)**
Granting a user the MODIFY, MONITOR, and OPERATE privileges on all virtual warehouses in Snowflake is equivalent to granting the global MANAGE WAREHOUSES privilege. These privileges collectively provide comprehensive control over virtual warehouses.
MODIFY Privilege:
Allows users to change the configuration of the virtual warehouse.
Includes resizing, suspending, and resuming the warehouse.
MONITOR Privilege:
Allows users to view the status and usage metrics of the virtual warehouse.
Enables monitoring of performance and workload.
OPERATE Privilege:
Grants the ability to start and stop the virtual warehouse.
Includes pausing and resuming operations as needed.
Snowflake Documentation: Warehouse Privileges

**NEW QUESTION: 120**
True or False: It is possible for a user to run a query against the query result cache without requiring an active Warehouse.
Freecram.net
**A.** True
**B.** False
**Answer: (SHOW ANSWER)**
Snowflake's architecture allows for the use of a query result cache that stores the results of queries for a period of time. If the same query is run again and the underlying data has not changed, Snowflake can retrieve the result from this cache without needing to re-run the query on an active warehouse, thus saving on compute resources.

**NEW QUESTION: 121**
What happens when an external or an internal stage is dropped? (Select TWO).
**A.** When dropping an external stage, the files are not removed and only the stage is dropped
**B.** When dropping an external stage, both the stage and the files within the stage are removed
**C.** When dropping an internal stage, the files are deleted with the stage and the files are recoverable
**D.** When dropping an internal stage, the files are deleted with the stage and the files are not recoverable
**E.** When dropping an internal stage, only selected files are deleted with the stage and are not recoverable
**Answer: (SHOW ANSWER)**
When an external stage is dropped in Snowflake, the reference to the external storage location is removed, but the actual files within the external storage (like Amazon S3, Google Cloud Storage, or Microsoft Azure) are not deleted. This means that the data remains intact in the external storage location, and only the stage object in Snowflake is removed.
On the other hand, when an internal stage is dropped, any files that were uploaded to the stage are deleted along with the stage itself. These files are not recoverable once the internal stage is dropped, as they are permanently removed from Snowflake's storage.
References:
[COF-C03] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Stages

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

**NEW QUESTION: 122**
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: A,D (LEAVE A REPLY)**
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: 123**
What can load data from an external stage right after the data is taken in?
**A.** COPY command scheduled with a task
**B.** Snowflake Kafka Connector
**C.** Auto-ingest Snowpipe
**D.** External table
**Answer: (SHOW ANSWER)**
The correct answer is C. Auto-ingest Snowpipe .
Auto-ingest Snowpipe continuously loads data from staged files as soon as notifications indicate that new files are available in cloud storage. It is designed for near-real-time ingestion from external stages.
Corrected wording:
A clearer version of the question is:
What can load data from an external stage shortly after the data arrives?
Why C is correct:
Auto-ingest Snowpipe uses cloud messaging notifications to detect when new files arrive in an external stage. Snowpipe then automatically loads the new files into a target Snowflake table using the COPY statement defined in the pipe.
Why the other options are incorrect:
A). A scheduled COPY command with a task can load data periodically, but it does not automatically load files right after arrival unless the schedule happens to align.
B). The Snowflake Kafka Connector streams data from Kafka topics into Snowflake, but the question specifically asks about loading from an external stage.
D). An external table lets Snowflake query data in external storage, but it does not load the data into Snowflake tables.
Official Snowflake documentation reference:
Snowflake documentation describes Snowpipe auto-ingest as a feature that loads data from staged files automatically after receiving event notifications from cloud storage.

Reference: Snowflake Documentation - Snowpipe; Snowflake Documentation - Auto-ingest Snowpipe; SnowPro Core Study Guide - Data Loading and Unloading.
==

**NEW QUESTION: 124**
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: 125**
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: 126**
Which Snowflake object will consume credits during automatic background maintenance?
**A.** View
**B.** Materialized view
**C.** External table
**D.** Table
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 127**
How can a user get the MOST detailed information about individual table storage details in Snowflake?
**A.** SHOW TABLES command
**B.** SHOW EXTERNAL TABLES command
**C.** TABLES view
**D.** TABLE STORAGE METRICS view
**Answer: (SHOW ANSWER)**
To get the most detailed information about individual table storage details in Snowflake, the TABLE STORAGE METRICS view should be used. This Information Schema view provides granular storage metrics for tables within Snowflake, including data related to the size of the table, the amount of data stored, and storage usage over time. It's an essential tool for administrators and users looking to monitor and optimize storage consumption and costs.
References:
Snowflake Documentation: Information Schema - TABLE STORAGE METRICS View

**NEW QUESTION: 128**
Which models are supported for Snowflake access control? (Select TWO)
**A.** UAC (User Access Control)
**B.** RBAC (Role-Based Access Control)
**C.** DAC (Discretionary Access Control)
**D.** OAC (Object Access Control)
**E.** EBAC (Entity-Based Access Control)
**Answer: (SHOW ANSWER)**

Snowflake primarily uses RBAC , where access is granted to roles, and roles are assigned to users. Snowflake also supports aspects of DAC , which allows privileges to be granted by users who have GRANT OPTION.
Reference:
Snowflake Docs: Understanding Access Control

**NEW QUESTION: 129**
What is one of the characteristics of data shares?
**A.** Data shares support full DML operations.
**B.** Data shares work by copying data to consumer accounts.
**C.** Data shares utilize secure views for sharing view objects.
**D.** Data shares are cloud agnostic and can cross regions by default.
**Answer: (SHOW ANSWER)**
Data sharing in Snowflake allows for live, read-only access to data across different Snowflake accounts without the need to copy or transfer the data. One of the characteristics of data shares is the ability to use secure views. Secure views are used within data shares to restrict the access of shared data, ensuring that consumers can only see the data that the provider intends to share, thereby preserving privacy and security.
References:
Snowflake Documentation: Understanding Secure Views in Data Sharing

**NEW QUESTION: 130**
How are clusters handled when the Economy scaling policy is set on a multi-cluster virtual warehouse?
**A.** A new cluster starts when the system detects that there is one or more queries queued waiting for an available cluster.
**B.** A new cluster will start every 60 seconds if there are queued queries that have been waiting for more than 30 seconds for an available cluster.
**C.** A new cluster starts only if the system estimates that the queued queries will require resources that will keep the new cluster busy for at least 6 minutes.
**D.** A cluster will shut down after 3 system checks indicating that the cluster has been under-utilized for a minimum of 6 minutes.
**Answer: (SHOW ANSWER)**
Under the Economy scaling policy, Snowflake delays starting a new cluster until it ' s sure the queries will utilize it effectively for ~6 minutes . This reduces costs but may slightly delay query execution.
Reference:
Snowflake Docs: Economy Scaling Policy

**NEW QUESTION: 131**
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: 132**
Which parameter can be set at the account level to set the minimum number of days for which Snowflake retains historical data in Time Travel?
**A.** DATA_RETENTION_TIME_IN_DAYS
**B.** MAX_DATA_EXTENSION_TIME_IN_DAYS
**C.** MIN_DATA_RETENTION_TIME_IN_DAYS
**D.** MAX CONCURRENCY LEVEL
**Answer: (SHOW ANSWER)**
The parameter DATA_RETENTION_TIME_IN_DAYS can be set at the account level to define the minimum number of days Snowflake retains historical data for Time Travel1.

**NEW QUESTION: 133**
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: 134**
Which function should be used to find the query ID of the second query executed in a current session?
**A.** Select LAST_QUERY_ID(-2)
**B.** Select LAST_QUERY_ID(2)
**C.** Select LAST_QUERY_ID(1)
**D.** Select LAST_QUERY_ID(2)
**Answer: (SHOW ANSWER)**
The correct function to find the query ID of the second query executed in the current session is SELECT LAST_QUERY_ID(-2). The LAST_QUERY_ID function returns the query ID for the most recent query executed in the session when called with no arguments. When used with an argument, it can retrieve the ID of previous queries within the same session, where -2 would reference the second most recent query executed.
References:

There's a clarification needed here; Snowflake's documentation indicates LAST_QUERY_ID() function does not accept arguments. It returns the ID of the last query executed in the session. To find the query ID of the second last executed query, users typically need to track query IDs manually or use session history views.

**NEW QUESTION: 135**
A team plans to build data-transformation pipelines using a Snowflake-supported development framework that provides compute pushdown and native integration with Snowflake features.
Which feature should the team use?
**A.** Snowpark API for Python
**B.** Snowflake JDBC driver
**C.** Snowflake ODBC driver
**D.** Snowflake Connector for Kafka
**Answer: (SHOW ANSWER)**
Snowpark provides language-specific APIs for building data-processing and transformation logic. Snowpark operations are translated into SQL and executed by Snowflake's compute engine. This provides compute pushdown because processing occurs inside Snowflake instead of requiring the complete dataset to be transferred to an external application.
The JDBC and ODBC drivers provide database connectivity. They allow applications to submit SQL statements and retrieve results, but they are not transformation frameworks comparable to Snowpark.
The Snowflake Connector for Kafka is primarily used to ingest streaming data from Apache Kafka into Snowflake.
Correction: The original option A said "Snowflake Connector for Python." That connector is a Python database driver.
The feature that provides DataFrame-based transformations and compute pushdown is the Snowpark API for Python.
References:
Snowflake Documentation - Snowpark Developer Guide for Python: https://docs.snowflake.com/en /developer-guide/snowpark/python/index
Snowflake Documentation - Snowpark API overview: https://docs.snowflake.com/en/developer-guide /snowpark/index
==

**NEW QUESTION: 136**
The data retention time for a schema is set to 5 days, and the data retention time for the database containing the schema is set to 10 days.
Based on these settings, what would be the default data retention time for a permanent table created in the schema?
**A.** 0 days
**B.** 1 day
**C.** 5 days
**D.** 10 days
**Answer: C (LEAVE A REPLY)**
The DATA_RETENTION_TIME_IN_DAYS parameter can be set at account, database, schema, or supported object levels. When a table does not explicitly set its own retention period, it inherits the applicable value from its nearest parent container. In this scenario, the database has a 10-day setting, but the schema inside that database has a more

specific 5-day setting. A permanent table created in that schema therefore inherits five days, not the database's ten days. An account-level MIN_DATA_RETENTION_TIME_IN_DAYS value could raise the effective retention period if it were configured, but no such setting is stated in the question. Because retention longer than one day requires Enterprise Edition or higher, the scenario also assumes an edition supporting extended Time Travel.
The correct answer is C.

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

**NEW QUESTION: 137**
Which grant will provide a user with access to data that is in a Snowflake share?
**A.** GRANT SELECT ON ALL TABLES IN DATABASE < SHARE > TO USER < USER > ;
**B.** GRANT READ ON ALL TABLES IN DATABASE < SHARE > TO ROLE < ROLE > ;
**C.** GRANT IMPORTED PRIVILEGES ON DATABASE < SHARE > TO ROLE < ROLE > ;
**D.** GRANT USE ON DATABASE < SHARE > TO ROLE < ROLE > ;
**Answer: (SHOW ANSWER)**
When a consumer imports a shared database , they must explicitly grant the IMPORTED PRIVILEGES to a role in order to access the shared objects.
Reference:
Snowflake Docs: Accessing Shared Data

**NEW QUESTION: 138**
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: 139**
In which scenarios would a user have to pay Cloud Services costs? (Select TWO).
**A.** Compute Credits = 50 Credits Cloud Services = 10
**B.** Compute Credits = 80 Credits Cloud Services = 5

**C.** Compute Credits = 10 Credits Cloud Services = 9
**D.** Compute Credits = 120 Credits Cloud Services = 10
**E.** Compute Credits = 200 Credits Cloud Services = 26
**Answer: (SHOW ANSWER)**
In Snowflake, Cloud Services costs are incurred when the Cloud Services usage exceeds 10% of the compute usage (measured in credits). Therefore, scenarios A and E would result in Cloud Services charges because the Cloud Services usage is more than 10% of the compute credits used.
References:
[COF-C03] SnowPro Core Certification Exam Study Guide
Snowflake's official documentation on billing and usage1

**NEW QUESTION: 140**
What is the minimum Snowflake edition required to use Dynamic Data Masking?
**A.** Standard
**B.** Enterprise
**C.** Business Critical
**D.** Virtual Private Snowflake (VPC)
**Answer: (SHOW ANSWER)**
The minimum Snowflake edition required to use Dynamic Data Masking is the Enterprise edition.This feature is not available in the Standard edition2.

**NEW QUESTION: 141**
When should a multi-cluster warehouse be used in auto-scaling mode?
**A.** When it is unknown how much compute power is needed
**B.** If the select statement contains a large number of temporary tables or Common Table Expressions (CTEs)
**C.** If the runtime of the executed query is very slow
**D.** When a large number of concurrent queries are run on the same warehouse
**Answer: (SHOW ANSWER)**
A multi-cluster warehouse should be used in auto-scaling mode when there is a need to handle a large number of concurrent queries.Auto-scaling allows Snowflake to automatically add or remove compute clusters to balance the load, ensuring that performance remains consistent during varying levels of demand

**NEW QUESTION: 142**
Which Snowflake tasks are managed in the cloud services layer? Select TWO.
**A.** Performance tuning
**B.** Query acceleration
**C.** Metadata management
**D.** Query processing
**E.** Query parsing
**Answer: (SHOW ANSWER)**
The correct answers are C. Metadata management and E. Query parsing .

Snowflake's architecture has three main layers:
Cloud Services Layer
Query Processing Layer / Compute Layer
Database Storage Layer
The cloud services layer coordinates Snowflake activities and handles services such as authentication, infrastructure management, metadata management, query parsing and optimization, access control, and transaction management.
Why C is correct:
Metadata management is handled by the cloud services layer. This includes metadata about tables, micro- partitions, query optimization statistics, and access-related metadata.
Why E is correct:
Query parsing and optimization are handled by the cloud services layer before execution is assigned to compute resources.
Why the other options are incorrect:
A). Performance tuning is not usually described as a discrete cloud services task. Snowflake automatically optimizes many operations, but this option is too broad.
B). Query acceleration uses additional compute resources and is a performance feature, not a core task of the cloud services layer.
D). Query processing is performed by virtual warehouses in the compute layer, not by the cloud services layer.
Official Snowflake documentation reference:
Snowflake documentation explains that the cloud services layer manages authentication, infrastructure management, metadata, query parsing and optimization, access control, and transaction management.
Reference: Snowflake Documentation - Snowflake architecture; Snowflake Documentation - Cloud services layer; SnowPro Core Study Guide - Snowflake Architecture.
==

**NEW QUESTION: 143**
How many network policies can be assigned to an account or specific user at a time?
**A.** One
**B.** Two
**C.** Three
**D.** Unlimited
**Answer: A (LEAVE A REPLY)**
According to my knowledge, a security administrator can create multiple network policies, but only one network policy can be active for an account or specific user at any given time. This ensures that there is a clear and consistent policy being applied without conflicts. References: Based on my internal knowledge as of
2021.

**NEW QUESTION: 144**
This query was executed:
Assuming the weather_events table has not changed, which query would benefit from cached results?
**A.** select distinct(severity) from weather_events; if run within 48 hours after the initial query execution

**B.** SELECT DISTINCT(severity) FROM weather_events we; if run within 24 hours after the initial query execution
**C.** select distinct (severity) from weather_events; if run within 24 hours after the initial query execution
**D.** SELECT DISTINCT(severity) FROM weather_events ; if run within 24 hours after the initial query execution
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 145**
Which table function will identify data that was loaded using COPY INTO <table> statements and also identify data loaded using Snowpipe?
**A.** DATA_TRANSFER_HISTORY
**B.** PIPE_USAGE_HISTORY
**C.** VALIDATE_PIPE_LOAD
**D.** COPY_HISTORY
**Answer: (SHOW ANSWER)**
The COPY_HISTORY table function in Snowflake is used to track data loading activities, including both manual COPY INTO <table> commands and automated data loading through Snowpipe. This function provides visibility into load operations, making it valuable for monitoring data ingestion, ensuring data integrity, and troubleshooting loading processes.

**NEW QUESTION: 146**
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

**NEW QUESTION: 147**
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 values3.

**NEW QUESTION: 148**
Which of the following statements describes a schema in Snowflake?

**A.** A logical grouping of objects that belongs to a single database
**B.** A logical grouping of objects that belongs to multiple databases
**C.** A named Snowflake object that includes all the information required to share a database
**D.** A uniquely identified Snowflake account within a business entity
**Answer: (SHOW ANSWER)**
A schema in Snowflake is a logical grouping of database objects, such as tables and views, that belongs to a single database.Each schema is part of a namespace in Snowflake, which is inferred from the current database and schema in use for the session5

**NEW QUESTION: 149**
Which URL provides access to files in Snowflake without authorization?
**A.** File URL
**B.** Scoped URL
**C.** Pre-signed URL
**D.** Scoped file URL
**Answer: (SHOW ANSWER)**
A Pre-signed URL provides access to files stored in Snowflake without requiring authorization at the time of access. This feature allows users to generate a URL with a limited validity period that grants temporary access to a file in a secure manner. It ' s particularly useful for sharing data with external parties or applications without the need for them to authenticate directly with Snowflake.
References:
Snowflake Documentation: Using Pre-signed URLs

**NEW QUESTION: 150**
What parameter setting will maximize parallel processing when unloading data from Snowflake?
**A.** SINGLE = TRUE
**B.** SINGLE = FALSE
**C.** DETAILED_OUTPUT
**D.** DETAILED OUTPUT = TRUE/FALSE
**Answer: (SHOW ANSWER)**
When SINGLE = FALSE is set during COPY INTO, Snowflake writes multiple output files , enabling parallelism . Setting SINGLE = TRUE forces one output file, reducing parallel processing.
Reference:
Snowflake Docs: COPY INTO Parameters

**NEW QUESTION: 151**
Which of the following describes how multiple Snowflake accounts in a single organization relate to various cloud providers?
**A.** Each Snowflake account can be hosted in a different cloud vendor and region.
**B.** Each Snowflake account must be hosted in a different cloud vendor and region
**C.** All Snowflake accounts must be hosted in the same cloud vendor and region

**D.** Each Snowflake account can be hosted in a different cloud vendor, but must be in the same region.
**Answer: (SHOW ANSWER)**
Snowflake's architecture allows for flexibility in account hosting across different cloud vendors and regions.
This means that within a single organization, different Snowflake accounts can be set up in various cloud environments, such as AWS, Azure, or GCP, and in different geographical regions. This allows organizations to leverage the global infrastructure of multiple cloud providers and optimize their data storage and computing needs based on regional requirements, data sovereignty laws, and other considerations.
https://docs.snowflake.com/en/user-guide/intro-regions.html

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

**NEW QUESTION: 152**
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.
==

**NEW QUESTION: 153**
How is Single Sign-On (SSO) authentication used in Snowflake?
**A.** SSO is an authentication method that allows a user to sign into multiple applications with a single set of credentials.
**B.** SSO is an integrated Snowflake feature powered by the Duo Security service which is managed completely by Snowflake.
**C.** SSO is an authentication method that uses a username and password in the API request header.
**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: 154**
What action is required to load data from a local file system into a Snowflake table?
**A.** Use the COPY INTO < table > command directly on the local file system.
**B.** Upload data files to a Snowflake external stage using the COPY INTO < table > and then use the GET command.
**C.** Use the PUT command directly on the local file system.
**D.** Upload data files to a Snowflake internal stage using the PUT command and then use the COPY INTO < table > command.
**Answer: (SHOW ANSWER)**
To load local files into Snowflake:
* First, use the PUT command to upload files from local storage to an internal stage .
* Then, use the COPY INTO < table > command to load data from the internal stage into the table.
Reference:
Snowflake Docs: Loading Data from Local Files

**NEW QUESTION: 155**
A user is configuring Snowflake CLI for multiple environments, such as dev, test, and prod. The user would like to avoid re-entering credentials every time and keep profiles organized. Which Snowflake CLI feature should be used to manage these environments efficiently?
**A.** Environment variables
**B.** Configuration file profiles
**C.** Authentication delegation
**D.** OAuth with a refresh token
**Answer: (SHOW ANSWER)**
Snowflake CLI supports configuration files with named connection profiles. These profiles allow users to organize different environments, such as dev, test, and prod, and reuse stored connection settings instead of re- entering them each time.
Reference: Snowflake Documentation - Snowflake CLI configuration; Connection profiles.

==

**NEW QUESTION: 156**
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: 157**
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 AP1 ON (NAME);
```

Which feature of this table will incur both compute and storage costs?
**A.** The row access policy
**B.** The Automatic Clustering
**C.** The Dynamic Data Masking policy
**D.** The not null constraint
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 158**
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: (SHOW ANSWER)**
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: 159**
What is required to set a stored procedure to EXECUTE AS? (Select TWO).
**A.** SCHEMA_USER
**B.** ACCOUNTADMIN role
**C.** USER
**D.** Owner ' s rights
**E.** Caller ' s rights
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 160**
A Snowflake user is creating a deployment pipeline for Snowflake objects.
What interface should be used to automate object deployment?
Freecram.net
**A.** ODBC driver
**B.** SnowSQL
**C.** SnowCD
**D.** Snowflake Horizon Catalog
**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 161**
How does an authentication policy affect security in a Snowflake account?
**A.** It controls how a client or user can access Snowflake.
**B.** It restricts inbound access to Snowflake data.
**C.** It restricts outbound traffic from Snowflake to external network destinations.
**D.** It provides users with the ability to use Single Sign-On, or SSO, to access Snowflake.
**Answer: (SHOW ANSWER)**
The correct answer is A. It controls how a client or user can access Snowflake .
An authentication policy controls authentication behavior for users or accounts. It determines which authentication methods are allowed and can be used to enforce security requirements for access to Snowflake.
Why A is correct:
Authentication policies control how users or clients authenticate to Snowflake. They can be used to allow or restrict authentication methods, helping organizations enforce secure access standards.
Why the other options are incorrect:
B). Restricting inbound access to Snowflake is handled by network policies, not authentication policies.
C). Restricting outbound traffic from Snowflake to external network destinations is handled by external access integrations and network rules, not authentication policies.
D). SSO can be part of an authentication strategy, but an authentication policy is broader than only enabling SSO.
Official Snowflake documentation reference:
Snowflake documentation describes authentication policies as security objects that control authentication methods for users in a Snowflake account.

Reference: Snowflake Documentation - Authentication policies; Snowflake Documentation - Network policies; SnowPro Core Study Guide - Security and Access Control.
==

**NEW QUESTION: 162**
What are the correct parameters for time travel and fail-safe in the Snowflake Enterprise Edition?
**A.** Default Time Travel Retention is set to 0 days.Maximum Time Travel Retention is 30 days.Fail Safe retention time is 1 day.
**B.** Default Time Travel Retention is set to 1 day.Maximum Time Travel Retention is 365 days.Fail Safe retention time is 7 days.
**C.** Default Time Travel Retention is set to 0 days.Maximum Time Travel Retention is 90 days.Fail Safe retention time is 7 days.
**D.** Default Time Travel Retention is set to 1 day.Maximum Time Travel Retention is 90 days.Fail Safe retention time is 7 days.
**E.** Default Time Travel Retention is set to 7 days.Maximum Time Travel Retention is 1 day.Fail Safe retention time is 90 days.
**F.** Default Time Travel Retention is set to 90 days.Maximum Time Travel Retention is 7 days.Fail Safe retention time is 356 days.
**Answer: (SHOW ANSWER)**
In the Snowflake Enterprise Edition, the default Time Travel retention is set to 1 day, the maximum Time Travel retention can be set up to 90 days, and the Fail-safe retention timeis 7 days3.

**NEW QUESTION: 163**
What virtual warehouse feature or setting will reduce the performance impact when running larger-than-average queries by offloading portions of the query processing work to shared compute resources?
**A.** Disabling auto-suspend
**B.** Using a multi-cluster virtual warehouse
**C.** Using the search optimization service
**D.** Using the query acceleration service
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 164**
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: (SHOW ANSWER)**
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: 165**
What will prevent unauthorized access to a Snowflake account from an unknown source?
**A.** Network policy
**B.** End-to-end encryption
**C.** Multi-Factor Authentication (MFA)
**D.** Role-Based Access Control (RBAC)
**Answer: (SHOW ANSWER)**
A network policy in Snowflake is used to restrict access to the Snowflake account from unauthorized or unknown sources.It allows administrators to specify allowed IP address ranges, thus preventing access from any IP addresses not listed in the policy1.

**NEW QUESTION: 166**
Which Snowflake object does not consume any storage costs?
**A.** Secure view
**B.** Materialized view
**C.** Temporary table
**D.** Transient table
**Answer: (SHOW ANSWER)**
Temporary tables do not consume any storage costs in Snowflake because they only exist for theduration of the session that created them and are automatically dropped when the session ends, thus incurring no long- term storage charges4. References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

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

**NEW QUESTION: 167**
Which column is used with streams to track changes to metadata?
**A.** METADATA$ROW_ID
**B.** METADATA$START_SCAN_TIME
**C.** METADATA$FILE_CONTENT_KEY
**D.** METADATA$FILE_ROW_NUMBER
**Answer: A (LEAVE A REPLY)**
The correct answer is A. METADATA$ROW_ID .

Snowflake streams record change data capture, or CDC, information for tables. When a stream tracks changes to a table, Snowflake adds metadata columns that describe the type of change and identify the affected row.
The standard stream metadata columns include:
METADATA$ACTION
METADATA$ISUPDATE
METADATA$ROW_ID
METADATA$ROW_ID is a unique and immutable identifier for a row. It is used to track the row involved in a change record.
Why A is correct:
METADATA$ROW_ID identifies the row associated with a change in a stream. This helps Snowflake track changed rows in stream output.
Why the other options are incorrect:
B). METADATA$START_SCAN_TIME is not a standard stream metadata column.
C). METADATA$FILE_CONTENT_KEY is file metadata used with staged files, not standard stream change tracking.
D). METADATA$FILE_ROW_NUMBER identifies a row number in a staged file during data loading, not a stream change row.
Official Snowflake documentation reference:
Snowflake documentation explains that streams expose metadata columns such as METADATA$ACTION, METADATA $ISUPDATE, and METADATA$ROW_ID. The METADATA$ROW_ID column uniquely identifies each row involved in a change.
Reference: Snowflake Documentation - Introduction to streams; Snowflake Documentation - Stream metadata columns; SnowPro Core Study Guide - SQL and Snowflake Objects.

**NEW QUESTION: 168**
Which Snowflake native tool can be used to diagnose and troubleshoot network connections?
**A.** SnowSQL
**B.** Snowflake Python connector
**C.** Snowsight
**D.** SnowCD
**Answer: (SHOW ANSWER)**
SnowSQL, Snowflake ' s command-line client, can be used to diagnose and troubleshoot network connections. SnowSQL provides various commands and options to test connectivity, configure network settings, and troubleshoot issues related to network connections between the client and Snowflake.
Snowflake Documentation: SnowSQL

**NEW QUESTION: 169**
Which queries can execute eligible portions in parallel by using the Query Acceleration Service? (Choose two.)
**A.** Large scans that delete many rows and columns.
**B.** Large scans that use an aggregation or a selective filter.
**C.** Queries with complex joins on small tables.
**D.** Small scans that read and write rows very slowly.

**E.** Large scans that insert many new rows.

**Answer: (SHOW ANSWER)**

The Query Acceleration Service uses serverless compute resources to execute eligible portions of a query plan in parallel with the virtual warehouse. Snowflake identifies two primary eligible patterns:
large scans containing an aggregation or highly selective filter, and large scans that insert or copy many new rows.
These patterns correspond directly to options B and E. Query acceleration is most useful for outlier queries with substantial scan work, unpredictable data volumes, or selective processing that can be divided across additional compute resources. Small scans generally do not contain enough parallelizable work to benefit. Complex joins on small tables are also not the primary eligibility pattern. Eligibility is determined at the query-plan level, so enabling the service does not guarantee that every submitted query will use it.

**NEW QUESTION: 170**

How many credits does a size 3X-Large virtual warehouse consume if it runs continuously for 2 hours?

**A.** 32

**B.** 64

**C.** 128

**D.** 256

**Answer: (SHOW ANSWER)**

In Snowflake, the consumption of credits by a virtual warehouse is determined by its size and the duration for which it runs. A size 3X-Large virtual warehouse consumes 128 credits if it runs continuously for 2 hours.
This consumption rate is based on the principle that larger warehouses, capable of providing greater computational resources and throughput, consume more credits per hour of operation. The specific rate of consumption is defined by Snowflake's pricing model and the scale of the virtual warehouse.
References: Snowflake Pricing Documentation

**NEW QUESTION: 171**

What task can a data consumer perform using third-party datasets from the Snowflake Marketplace?

**A.** Republish listings from the data provider.

**B.** Integrate Marketplace data with data in the consumer's Snowflake environment.

**C.** Modify the data structures and schemas in the provider's dataset.

**D.** Export Marketplace datasets to local files for offline analysis.

**Answer: (SHOW ANSWER)**

Snowflake Marketplace allows consumers to discover and access third-party datasets without building separate ingestion pipelines for each provider. Consumers can combine Marketplace data with data already stored in their Snowflake environment and query the combined information to produce new analytics and business insights. The provider continues to control and update the shared data product, while the consumer generally receives read-only access through a listing. Consumers cannot modify the provider's underlying schemas or source data, and imported shared databases cannot be re-shared as the consumer's own listing. Exporting data to local files is not the defining Marketplace capability described in this question. Therefore, integrating third-party Marketplace data with the consumer's existing Snowflake data is the supported task represented by option B.

**NEW QUESTION: 172**
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: 173**
Which option will replace NULL values with an empty string while unloading data into a file?
**A.** FIELD_OPTIONALLY_ENCLOSED_BY = NONE
**B.** null_if=( ' NULL ' )
**C.** EMPTY_FIELD_AS_NULL = TRUE
**D.** null_if= ( ' ' )
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 174**
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: 175**
A user is automating a data pipeline using Snowflake CLI and needs to pass variables into SQL scripts.
Which variable-templating syntaxes are supported? (Choose three.)
**A.** STANDARD
**B.** LEGACY
**C.** JINJA
**D.** NONE
**E.** ALL
**Answer: (SHOW ANSWER)**
Snowflake CLI supports three actual variable-templating syntaxes for SQL execution: STANDARD, LEGACY, and JINJA. STANDARD uses the native Snowflake CLI form, such as < % variable_name
% > . LEGACY supports SnowSQL-style variables, such as & variable_name or & {variable_name}.
JINJA uses double-brace expressions, such as {{ variable_name }}. ALL and NONE are not separate templating syntaxes. They are control values: ALL enables every supported syntax, while NONE disables template resolution. Snowflake currently enables STANDARD and LEGACY by default, whereas JINJA must normally be enabled explicitly. Therefore, the original "Choose two" instruction and the supplied B, C answer are inconsistent with the official documentation. The corrected question must say "Choose three," and the verified answer is A, B, C.

**NEW QUESTION: 176**
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: C,E (LEAVE A REPLY)**
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: 177**
What best practice recommendations will help prevent timeouts when using the PUT command to load large data sets? (Select TWO).
**A.** Compress the files before loading.
**B.** Use a semi-structured file format.
**C.** Increase the PARALLEL option value.
**D.** Load the data into a table stage.
**E.** Enable the overwrite option.
**Answer: A,C (LEAVE A REPLY)**
To avoid timeouts during large data uploads with the PUT command in Snowflake, it is recommended to:
Compress files before loading: Compressed files are smaller and upload faster, reducing the risk of timeouts.
Increase the PARALLEL option value: This option allows more simultaneous upload threads, improving upload speed and efficiency for large datasets.
Semi-structured file formats and table staging do not directly impact timeouts, while enabling overwrite does not prevent timeouts but rather controls overwriting of existing files.

**NEW QUESTION: 178**
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 isparticularly 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: 179**
How do secure views compare to non-secure views in Snowflake?
**A.** Secure views execute slowly compared to non-secure views.
**B.** Non-secure views are preferred over secure views when sharing data.
**C.** Secure views are similar to materialized views in that they are the most performant.
**D.** There are no performance differences between secure and non-secure views.
**Answer: (SHOW ANSWER)**
Secure views and non-secure views in Snowflake are differentiated primarily by their handling of data access and security rather than performance characteristics. A secure view enforces row-level security and ensures that the view definition is hidden from the users. However, in terms of performance, secure views do not inherently execute slower or faster than non-secure views. The performance of both types of views depends more on other factors such as

underlying table design, query complexity, and system workload rather than the security features embedded in the views themselves.
References:
Snowflake Documentation on Views: This section provides an overview of both secure and non-secure views, clarifying that the main difference lies in security features rather than performance, thus supporting the assertion that there are no inherent performance differences.

**NEW QUESTION: 180**
Which components make up the Snowflake Cortex AI suite? Select TWO.
**A.** Document AI
**B.** Snowpipe
**C.** Streams
**D.** Secure views
**E.** Cortex Analyst
**Answer: (SHOW ANSWER)**
The correct answers are A. Document AI and E. Cortex Analyst .
Snowflake Cortex is Snowflake's integrated AI and machine learning capability set. It includes features for working with language models, document extraction, natural language analytics, and AI-powered application development.
Why A is correct:
Document AI is part of Snowflake Cortex. It is used to extract information from unstructured documents using AI.
Why E is correct:
Cortex Analyst is part of Snowflake Cortex. It enables natural language questions over structured data by using semantic models and Cortex capabilities.
Why the other options are incorrect:
B). Snowpipe is a data ingestion service, not part of the Cortex AI suite.
C). Streams are change data capture objects used to track table changes, not Cortex AI components.
D). Secure views are security/governance objects used to protect view definitions and support secure data sharing, not Cortex AI components.
Official Snowflake documentation reference:
Snowflake documentation describes Snowflake Cortex as a set of AI and ML capabilities and includes features such as Document AI and Cortex Analyst. Snowpipe, streams, and secure views belong to other Snowflake feature areas.
Reference: Snowflake Documentation - Snowflake Cortex; Snowflake Documentation - Document AI; Snowflake Documentation - Cortex Analyst; SnowPro Core Study Guide - Snowflake Architecture.

**NEW QUESTION: 181**
What languages can be used to write User-Defined Functions, or UDFs? Select TWO.
**A.** Groovy
**B.** JavaScript
**C.** Python
**D.** Spark
**E.** Golang

**Answer: (SHOW ANSWER)**

The correct answers are B. JavaScript and C. Python .

Snowflake supports User-Defined Functions, or UDFs, in multiple languages. Common supported languages include SQL, JavaScript, Java, Python, and Scala. From the options provided, the supported UDF languages are JavaScript and Python.

Why B is correct:

JavaScript is supported for Snowflake UDFs. JavaScript UDFs allow users to write procedural logic that can be called from SQL.

Why C is correct:

Python is supported for Snowflake UDFs. Python UDFs allow users to use Python code inside Snowflake for custom processing.

Why the other options are incorrect:

A). Groovy is not a supported Snowflake UDF language.

D). Spark is a data processing engine, not a Snowflake UDF programming language.

E). Golang is not a supported Snowflake UDF language.

Official Snowflake documentation reference:

Snowflake documentation lists supported UDF languages, including SQL, JavaScript, Java, Python, and Scala.

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

==

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

**NEW QUESTION: 182**

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

**A.** ROW_NUMBER

**B.** TABLE

**C.** TABULAR

**D.** VALUES

**Answer: (SHOW ANSWER)**

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

**NEW QUESTION: 183**

The Query Profile shows the metrics of a query that ran on a size Large virtual warehouse.
What steps will improve the query performance? (Select TWO).
**A.** Increase the WAREHOUSE_SIZE.
**B.** Run the warehouse in Maximized mode.
**C.** Update the query WHERE clause.
**D.** Increase the MAX_CLUSTER_COUNT.
**E.** Increase the warehouse AUTO_SUSPEND duration.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 184**
Which activities are managed by Slowflake ' s Cloud Services layer? (Select TWO).
**A.** Authorisation
**B.** Access delegation
**C.** Data pruning
**D.** Data compression
**E.** Query parsing and optimization
**Answer: (SHOW ANSWER)**
Snowflake ' s Cloud Services layer is responsible for managing various aspects of the platform that are not directly related to computing or storage. Specifically, it handles authorisation, ensuring that users have appropriate access rights to perform actions or access data. Additionally, it takes care of query parsing and optimization, interpreting SQL queries and optimizing their execution plans for better performance. This layer abstracts much of the platform ' s complexity, allowing users to focus on their data and queries without managing the underlying infrastructure.
References: Snowflake Architecture Documentation

**NEW QUESTION: 185**
What is the data lineage feature used for in Snowflake?
**A.** To track how records are read by a user
**B.** To track dependencies at the object level
**C.** To track how data has been transformed among columns
**D.** To track how transformation processes have impacted individual records
**Answer: (SHOW ANSWER)**
The correct answer is C. To track how data has been transformed among columns .
Snowflake's data lineage capability is used to understand how data flows and changes across objects. It helps users see relationships between source and target objects and understand how data is transformed, especially at the column level.
Corrected wording:
The original option said "records," but Snowflake data lineage is more accurately described as tracking how data is transformed among columns, not individual row-level record transformations.
Why C is correct:

Data lineage helps track how data moves from source columns to target columns through operations such as views, CTAS statements, INSERT statements, and other transformations. This is useful for governance, impact analysis, auditing, and understanding downstream dependencies.
Why the other options are incorrect:
A). Tracking how records are read by a user is more closely related to access history or query history, not data lineage.
B). Tracking dependencies at the object level is related to object dependencies. Data lineage is more specific because it can show how data flows between columns and objects.
D). Snowflake data lineage is not primarily a row-level feature for tracking the impact of transformations on individual records.
Official Snowflake documentation reference:
Snowflake documentation describes data lineage as a governance capability that helps track how data moves and transforms across Snowflake objects, including column-level relationships.
Reference: Snowflake Documentation - Data lineage; Snowflake Documentation - Object dependencies; SnowPro Core Study Guide - Data Protection and Governance.
==

**NEW QUESTION: 186**
How can a transient table be converted to a permanent table?
**A.** Use the ALTER TABLE ... RENAME command.
**B.** Use the CREATE TABLE AS SELECT, or CTAS, command.
**C.** Use the ALTER TABLE ... SWAP WITH command.
**D.** Use the CREATE TABLE ... CLONE command.
**Answer: (SHOW ANSWER)**
The correct answer is B. Use the CREATE TABLE AS SELECT, or CTAS, command .
A transient table cannot be directly converted into a permanent table by changing a table property. To create a permanent version of the data, create a new permanent table from the transient table data.
Why B is correct:
A CREATE TABLE AS SELECT statement can create a new permanent table from the data in the transient table. Since permanent is the default table type unless otherwise specified, omitting TRANSIENT creates a permanent table.
Example:
CREATE TABLE permanent_table AS
SELECT *
FROM transient_table;
Why the other options are incorrect:
A). ALTER TABLE ... RENAME only changes the table name. It does not change table type.
C). ALTER TABLE ... SWAP WITH swaps table names and metadata between two existing tables, but it is not the method to convert a transient table into a permanent table.
D). CREATE TABLE ... CLONE preserves the source table type behavior in this context and is not the recommended answer for converting transient data into a permanent table.
Official Snowflake documentation reference:

Snowflake documentation explains that the table type cannot be changed after table creation. To change from transient to permanent, create a new permanent table and copy the data, such as with CTAS.
Reference: Snowflake Documentation - Transient tables; Snowflake Documentation - CREATE TABLE AS SELECT; SnowPro Core Study Guide - SQL and Snowflake Objects.
==

**NEW QUESTION: 187**
How does Snowflake Fail-safe protect data in a permanent table?
**A.** Fail-safe makes data available up to 1 day, recoverable by user operations.
**B.** Fail-safe makes data available for 7 days, recoverable by user operations.
**C.** Fail-safe makes data available for 7 days, recoverable only by Snowflake Support.
**D.** Fail-safe makes data available up to 1 day, recoverable only by Snowflake Support.
**Answer: (SHOW ANSWER)**
Snowflake's Fail-safe provides a 7-day period during which data in a permanent table may be recoverable, but only by Snowflake Support, not by user operations3.

**NEW QUESTION: 188**
Which activities are included in the Cloud Sen/ices 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 includes activities such as user authentication and infrastructure management. This layer coordinates activities across Snowflake, including security enforcement, query compilation and optimization, and more

**NEW QUESTION: 189**
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: 190**

A table is created with a column set to a data type of CHARACTER(25). If the table is viewed in the database section of Snowsight, what data type will be shown for the column?

Freecram.net

**A.** STRING(25)
**B.** TEXT(25)
**C.** VARCHAR(25)
**D.** CHARACTER(25)
**Answer: (SHOW ANSWER)**
The correct answer is C. VARCHAR(25) .
In Snowflake, CHAR, CHARACTER, STRING, and TEXT are synonyms for VARCHAR. Snowflake stores and displays these character string data types as VARCHAR.
Corrected typo:
Option D was written as CHARACTERS). The corrected option is CHARACTER(25) .
Why C is correct:
A column created as CHARACTER(25) is treated as VARCHAR(25) in Snowflake metadata and display output.
Why the other options are incorrect:
A). STRING is a synonym for VARCHAR, but Snowflake commonly displays the canonical type as VARCHAR.
B). TEXT is also a synonym for VARCHAR, but it is not the canonical display type.
D). CHARACTER is accepted when defining the column, but Snowflake displays it as VARCHAR.
Official Snowflake documentation reference:
Snowflake documentation explains that string data type synonyms such as CHAR, CHARACTER, STRING, and TEXT are equivalent to VARCHAR.
Reference: Snowflake Documentation - String and binary data types; SnowPro Core Study Guide - SQL and Snowflake Objects.
==

**NEW QUESTION: 191**
Which Snowflake layer is associated with virtual warehouses?
**A.** Cloud services
**B.** Query processing
**C.** Elastic memory
**D.** Database storage
**Answer: B (LEAVE A REPLY)**
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: 192**
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: 193**
During periods of warehouse contention which parameter controls the maximum length of time a warehouse will hold a query for processing?
**A.** STATEMENT_TIMEOUT__IN__SECONDS
**B.** STATEMENT_QUEUED_TIMEOUT_IN_SECONDS
**C.** MAX_CONCURRENCY__LEVEL
**D.** QUERY_TIMEOUT_IN_SECONDS

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

**NEW QUESTION: 194**
The Snowflake Search Optimization Services supports improved performance of which kind of query?
**A.** Queries against large tables where frequent DML occurs
**B.** Queries against tables larger than 1 TB
**C.** Selective point lookup queries
**D.** Queries against a subset of columns in a table
**Answer: (SHOW ANSWER)**
The Snowflake Search Optimization Service is designed to support improved performance for selective point lookup queries.These are queries that retrieve specific records from a database, often based on a unique identifier or a small set of criteria3.

**NEW QUESTION: 195**
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: 196**
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 ExamDiscuss.com for Helping Passing COF-C03 Exam! ExamDiscuss.com now offer the **newest COF-C03 exam dumps**, the ExamDiscuss.com COF-C03 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com COF-C03 dumps with Test Engine here: https://www.examdiscuss.com/Snowflake/exam/COF-C03/premium/ (**1187** Q&As Dumps, **35%OFF** **Special Discount Code:** **freecram**)

**NEW QUESTION: 197**
Assume there is a table consisting of five micro-partitions with values ranging from A to Z.
Which diagram indicates a well-clustered table?
**A.** All letters A-Z are randomly mixed across every micro-partition, each partition has a full wide range of values
**B.** Alphabetical values are scattered randomly across partitions with heavy overlapping ranges between micro-partitions
**C.** The A-Z range is split into 5 narrow, continuous alphabet segments (A-F, G-L, M-R, S-W, X-Z), each partition holds one tight range with almost no overlap
**D.** One single micro-partition stores nearly all data from A-Z; the other four partitions contain very few random rows
**Answer: (SHOW ANSWER)**
A well-clustered table in Snowflake means that the data is organized in such a way that related data points are stored close to each other within the micro-partitions. This optimizes query performance by reducing the amount of scanned data. The diagram indicated by option C shows a well-clustered table, as it likely represents a more evenly distributed range of values across the micro-partitions1.
References = Snowflake Micro-partitions & TableClustering

**NEW QUESTION: 198**
A data provider offers a large financial dataset through Snowflake Marketplace. The provider wants to share the dataset with specific investment firms that already use Snowflake.
The provider does not want the listing to appear in Marketplace search results, and only invited customers should be able to view and subscribe to it.
Which type of Snowflake listing should the provider create?
**A.** Public listing
**B.** Private listing
**C.** Personalized Data Exchange

**D.** Provider profile listing

**Answer: (SHOW ANSWER)**

A private listing is made available only to selected consumer accounts. It is not generally discoverable by all Snowflake Marketplace users, and the provider controls which accounts can access it.

A public listing is discoverable through Snowflake Marketplace and is intended for broad availability. A provider profile describes the provider and its offerings but is not itself the mechanism for privately sharing a dataset.

"Personalized Data Exchange" is not the correct listing type for this requirement. Snowflake previously used the term Data Exchange for private data-sharing hubs, but private listings are the appropriate mechanism when a provider wants to offer a listing directly to selected consumers.

References:

Snowflake Documentation - About listings: https://docs.snowflake.com/en/collaboration/provider-listings- about

Snowflake Documentation - Creating and publishing listings: https://docs.snowflake.com/en/collaboration /provider-listings-creating-publishing

==

**NEW QUESTION: 199**

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

**A.** SELECT

**B.** MERGE

**C.** INSERT

**D.** DELETE

**E.** UPDATE

**Answer: B,C (LEAVE A REPLY)**

**NEW QUESTION: 200**

Which metrics are provided in the Query Profile?

**A.** Query execution time and number of rows processed

**B.** Data retention period and backup schedules

**C.** Roles and permission of users running the queries

**D.** The costs of running queries by user

**Answer: (SHOW ANSWER)**

The Query Profile tool provides a visual breakdown of query execution including:

* Execution time per step

* Rows processed , scanned, or returned This helps analyze performance bottlenecks and optimize SQL.

Reference:

Snowflake Docs: Query Profile

**NEW QUESTION: 201**

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: 202**
What is a key feature of Snowflake architecture?
**A.** Zero-copy cloning creates a mirror copy of a database that updates with the original
**B.** Software updates are automatically applied on a quarterly basis
**C.** Snowflake eliminates resource contention with its virtual warehouse implementation
**D.** Multi-cluster warehouses allow users to run a query that spans across multiple clusters
**E.** Snowflake automatically sorts DATE columns during ingest for fast retrieval by date
**Answer: (SHOW ANSWER)**
One of the key features of Snowflake's architecture is its unique approach to eliminating resource contention through the use of virtual warehouses. This is achieved by separating storage and compute resources, allowing multiple virtual warehouses to operate independently on the same data without affecting each other. This means that different workloads, such as loading data, running queries, or performing complex analytics, can be processed simultaneously without any performance degradation due to resource contention.
References:
Snowflake Documentation on Virtual Warehouses
SnowPro Core Certification Study Guide

**NEW QUESTION: 203**
What is the LEAST-PRIVILEGED role that can enable replication for both source and target accounts by default?
**A.** ACCOUNTADMIN
**B.** ORGADMIN
**C.** SECURITYADMIN
**D.** SYSADMIN
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 204**
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: 205**
A running virtual warehouse is suspended.
What is the MINIMUM amount of time that the warehouse will incur charges for when it is restarted?
**A.** 1 second
**B.** 60 seconds
**C.** 5 minutes
**D.** 60 minutes
**Answer: (SHOW ANSWER)**
When a running virtual warehouse in Snowflake is suspended and then restarted, the minimum amount of time it will incur charges for is 60 seconds2.

**NEW QUESTION: 206**
How can storage costs be MINIMIZED while maintaining the history and integrity of a large, permanent table that is updated frequently?
**A.** Enable automatic clustering
**B.** Convert the table type to transient.
**C.** Decrease the Time Travel retention period
**D.** Use zero-copy clones of the table.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 207**
A data team needs to connect Snowflake to a remote Git repository over a public network using a personal access token for authentication.
Which components should be created in Snowflake to implement this connection? (Choose two.)
**A.** A storage integration
**B.** Azure Private Link
**C.** A secret
**D.** An API integration

**E.** A security integration

**Answer: (SHOW ANSWER)**

For a Git repository reached over the public network with a personal access token, Snowflake requires a secret and an API integration. The secret stores the repository username and token, typically as a PASSWORD-type secret. The API integration defines the Git HTTPS API provider, allowed repository URL prefixes, and the authentication secrets Snowflake is permitted to use. A Git repository object can then reference both components through API_INTEGRATION and GIT_CREDENTIALS. A storage integration is used for external cloud storage, not Git authentication. Azure Private Link applies to private connectivity, and a security integration is not the object used for this Git HTTPS configuration. Therefore, the answer supplied in the candidate file has been corrected from C and E to C and D.

**NEW QUESTION: 208**

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: 209**

A Snowflake table that is loaded using a Kafka connector has a schema consisting of which two variant columns? (Select TWO).

**A.** RECORD_TIMESTAMP

**B.** RECORD_CONTENT

**C.** RECORDKEY

**D.** RECORD_SESSION

**E.** RECORD_METADATA

**Answer: (SHOW ANSWER)**

When using theSnowflake Kafka connector, the table schema includes two importantvariant columns:
RECORD_TIMESTAMP: This column stores the timestamp from the Kafka record, enabling time-based analysis of incoming data.
RECORDKEY: This captures the unique key of each Kafka message, useful for uniquely identifying records or managing deduplication.
These columns ensure that each message's metadata and key information are preserved, facilitating data analysis and real-time processing tasks in Snowflake.

**NEW QUESTION: 210**

What is the MINIMUM role required to set the value for the parameter ENABLE_ACCOUNT_DATABASE_REPLICATION?

freedram.net

**A.** ACCOUNTADMIN

**B.** SECURITYADMIN
**C.** SYSADMIN
**D.** ORGADMIN
**Answer: (SHOW ANSWER)**
TheENABLE_ACCOUNT_DATABASE_REPLICATIONparameter is a critical setting in Snowflake that allows or restricts the replication of databases across Snowflake accounts. Given the significant impact of this parameter on data management and security, only theACCOUNTADMINrole has the minimum required privileges to set or modify it. This ensures that only users with the highest level of access and responsibility within the Snowflake environment can control database replication settings, maintaining strict governance and security standards.
References: Snowflake Documentation on Database Replication

**NEW QUESTION: 211**
Snowflake users can create a resource monitor at which levels? (Select TWO).
**A.** User level
**B.** Pipe level
**C.** Account level
**D.** Cloud services level
**E.** Virtual warehouse level
**Answer: (SHOW ANSWER)**
Resource monitors in Snowflake are tools used to track and control the consumption of compute resources, ensuring that usage stays within defined limits. These monitors can be created at the account level, allowing administrators to set overall resource consumption limits for the entire Snowflake account. Additionally, resource monitors can be set at the virtual warehouse level, enabling more granular control over the resources consumed by individual warehouses. This dual-level capability allows organizations to manage their Snowflake usage efficiently, preventing unexpected costs and optimizing performance.
References:Snowflake Documentation on Resource Monitors
Top of Form

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

**NEW QUESTION: 212**
What can be interpreted based on the number of nodes reported in a Query Profile?
**A.** Query execution time
**B.** Compute required for metadata management
**C.** Network bandwidth during query execution

**D.** Data spillage
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 213**
A query against which table will provide details about bytes consumed by active, Time Travel, and Fail-safe storage?
**A.** METERING_HISTORY
**B.** TABLE_STORAGE_METRICS
**C.** CREDIT_USAGE_HISTORY
**D.** DATA_TRANSFER_HISTORY
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 214**
Which statement describes Snowflake tables?
**A.** Snowflake tables arc 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 lo perform optimally.
**D.** Snowflake tables are owned by a use.
**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: 215**
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: 216**
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: 217**
A JSON file, that contains lots of dates and arrays, needs to be processed in Snowflake. The user wants to ensure optimal performance while querying the data.
How can this be achieved?
**A.** Flatten the data and store it in structured data types in a flattened table. Query the table.
**B.** Store the data in a table with a variant data type. Query the table.
**C.** Store the data in a table with a vai : ant data type and include STRIP_NULL_VALUES while loading the table. Query the table.
**D.** Store the data in an external stage and create views on top of it. Query the views.
**Answer: (SHOW ANSWER)**
Storing JSON data in a table with a VARIANT data type is optimal for querying because it allows Snowflake to leverage its semi-structured data capabilities. This approach enables efficient storage and querying without the need for flattening the data, which can be performance-intensive1.

**NEW QUESTION: 218**
Which function will generate a URL that can be accessed by a non-Snowflake user?
**A.** BUILD_SCOPED_FILE_URL
**B.** GET_PRESIGNED_URL
**C.** GET_STAGE_LOCATION
**D.** BUILD_STAGE_FILE_URL
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 219**
Which Snowflake feature or service is primarily used for managing and monitoring data and user activities?
**A.** Streamlit
**B.** SnowSQL

**C.** Snowsight
**D.** Snowflake Marketplace
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 220**
Which languages requite 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 UDF2.

**NEW QUESTION: 221**
The first query run on a virtual warehouse takes longer to complete than the next execution of the same query one hour later.
What causes this behavior?
**A.** The network connection is slower because of high traffic.
**B.** The warehouse copies all data from storage before executing the first query.
**C.** The first query is compiled by the database, and the second query uses a permanently precompiled version.
**D.** The warehouse was initially suspended, and the second query used cached data.
**Answer: (SHOW ANSWER)**
A warehouse's local table-data cache is dropped when the warehouse suspends. The first query submitted after suspension must resume the warehouse and read the required data without the benefit of a warm warehouse cache, so it can take longer. If the warehouse remains active, later queries can reuse cached table data. An identical query can also reuse a persisted query result when Snowflake's reuse conditions are satisfied, including unchanged source data, identical query text, valid privileges, and an available cached result. Among the supplied choices, option D best explains the observed behavior: the warehouse was initially suspended, while the later execution benefited from cached data. Network traffic and permanently precompiled query versions are not the primary explanations documented by Snowflake.

**NEW QUESTION: 222**
Using which object-level parameters will help limit query processing and concurrency slowdowns? (Select TWO).
**A.** QUERY_TAG
**B.** STATEMENT_TIMEOUT_IN_SECONDS
**C.** MAX_CONCURRENCY_LEVEL
**D.** AUTO_SUSPEND
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 223**
What do these Query Profile metrics indicate about the query micro-partition pruning?
**A.** No micro-partitions were pruned during execution.
**B.** Since the bytes scanned is less than the bytes written to results some pruning has occurred.
**C.** The size of the data queried was too small to trigger micro-partition pruning.
**D.** The query used the local disk cache and pruning was not required.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 224**
What does a masking policy consist of in Snowflake?
**A.** A single data type, with one or more conditions, and one or more masking functions
**B.** A single data type, with only one condition, and only one masking function
**C.** Multiple data types, with only one condition, and one or more masking functions
**D.** Multiple data types, with one or more conditions, and one or more masking functions
**Answer: (SHOW ANSWER)**
A masking policy in Snowflake consists of a single data type, with one or more conditions, and one or more masking functions.These components define how the data is masked based on the specified conditions3.

**NEW QUESTION: 225**
What is the purpose of this Snowflake function? Select TWO.
SYSTEM$PIPE_STATUS( ' < pipe_name > ' )
**A.** To identify the name of a pipe
**B.** To modify the configuration of a pipe
**C.** To determine if the pipe is suspended
**D.** To troubleshoot errors in a pipe execution
**E.** To list all staged files that include data unloaded from a pipe
**Answer: (SHOW ANSWER)**
The correct answers are C. To determine if the pipe is suspended and D. To troubleshoot errors in a pipe execution .
SYSTEM$PIPE_STATUS returns status information about a Snowpipe pipe. This function is useful for monitoring Snowpipe health and troubleshooting ingestion issues.
Why C is correct:
The output of SYSTEM$PIPE_STATUS includes information that can indicate whether a pipe is paused or running, helping determine whether ingestion is suspended.
Why D is correct:
The function can return diagnostic information about pipe execution, including pending files, notification channel status, and error-related information. This helps troubleshoot Snowpipe loading problems.
Example:
SELECT SYSTEM$PIPE_STATUS( ' my_pipe ' );
Why the other options are incorrect:
A). The pipe name is supplied as input to the function; the function is not primarily used to identify the pipe name.

B). SYSTEM$PIPE_STATUS is informational only. It does not modify pipe configuration. Configuration changes are made with commands such as ALTER PIPE.
E). Snowpipe loads data into tables. It does not unload data. Listing staged files is done using stage listing commands or directory tables, not SYSTEM$PIPE_STATUS.
Official Snowflake documentation reference:
Snowflake documentation describes SYSTEM$PIPE_STATUS as a system function that retrieves a JSON representation of the current status of a pipe, useful for monitoring and troubleshooting Snowpipe.
Reference: Snowflake Documentation - SYSTEM$PIPE_STATUS; Snowflake Documentation - Snowpipe monitoring and troubleshooting; SnowPro Core Study Guide - Data Loading and Unloading.
==

**NEW QUESTION: 226**
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 Snowflakeuser can use the VALIDATION_MODE parameter within the SQL statement to test the files for errors without loading them3.

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

**NEW QUESTION: 227**
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: 228**
Which of the following are characteristics of Snowflake virtual warehouses? (Choose two.)
**A.** Auto-resume applies only to the last warehouse that was started in a multi-cluster warehouse.
**B.** The ability to auto-suspend a warehouse is only available in the Enterprise edition or above.
**C.** SnowSQL supports both a configuration file and a command line option for specifying a default warehouse.
**D.** A user cannot specify a default warehouse when using the ODBC driver.
**E.** The default virtual warehouse size can be changed at any time.
**Answer: (SHOW ANSWER)**
Snowflake virtual warehouses support a configuration file and command line options in SnowSQL to specify a default warehouse, which is characteristic C. Additionally, the size of a virtual warehouse can be changed at any time, which is characteristic E. Thesefeatures provide flexibility and ease of use in managing compute resources2.
References = [COF-C03] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

**NEW QUESTION: 229**
An engineer has a connection named dev_project in the Snowflake CLI config.toml file. The connection specifies a role other than SYSADMIN.
The engineer then runs:
snow sql -c dev_project
When the engineer executes:
SELECT CURRENT_ROLE();
the query returns SYSADMIN.
Why did Snowflake CLI use the SYSADMIN role?
**A.** SYSADMIN is the engineer's default role and cannot be overridden by the config.toml file.
**B.** The [connections.dev_project] connection in the config.toml file was missing a warehouse parameter, which caused Snowflake CLI to ignore the role parameter.
**C.** A SNOWFLAKE_ROLE environment variable was set to SYSADMIN, and environment variables take precedence over values in the config.toml file.
**D.** The config.toml file is used only when the connection and role are also supplied as command-line parameters.
**Answer: (SHOW ANSWER)**
Snowflake CLI can obtain connection parameters from multiple sources. When the same connection parameter is specified in more than one place, the higher-precedence source overrides the lower-precedence source.
An environment variable such as SNOWFLAKE_ROLE overrides the role configured for the connection in config.toml. Therefore, when SNOWFLAKE_ROLE=SYSADMIN is present, the CLI session uses SYSADMIN, even when the selected connection specifies another role.
A missing warehouse does not cause Snowflake CLI to ignore the configured role. A user's default role can also be overridden when another valid role is explicitly supplied through a higher-precedence configuration source.
References:
Snowflake Documentation - Configuring connections for Snowflake CLI: https://docs.snowflake.com/en/developer-guide/snowflake-cli/connecting/configure-connections
Snowflake Documentation - Connecting to Snowflake with Snowflake CLI: https://docs.snowflake.com/en

/developer-guide/snowflake-cli/connecting/connect
==

**NEW QUESTION: 230**
Which Snowflake objects consume storage? (Select TWO)
**A.** External tables
**B.** Views
**C.** Temporary schemas
**D.** Materialized views
**E.** Permanent tables
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 231**
A company strongly encourages all Snowflake users to self-enroll in Snowflake's default Multi-Factor Authentication (MFA) service to provide increased login security for users connecting to Snowflake.
Which application will the Snowflake users need to install on their devices in order to connect with MFA?
**A.** Okta Verify
**B.** Duo Mobile
**C.** Microsoft Authenticator
**D.** Google Authenticator
**Answer: (SHOW ANSWER)**
Snowflake's default Multi-Factor Authentication (MFA) service is powered by Duo Security. Users are required to install the Duo Mobile application on their devices to use MFA for increased login security when connecting to Snowflake.This service is managed entirely by Snowflake, and users do not need to sign up separately with Duo1.

**NEW QUESTION: 232**
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:

Snowflake Documentation onSemi-Structured Data Types
SnowPro Core Certification Study Guide

**NEW QUESTION: 233**
When referring to User-Defined Function (UDF) names in Snowflake, what does the term overloading mean?
**A.** There are multiple SOL 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: 234**
What is the Fail-safe retention period for transient and temporary tables?
**A.** 0 days
**B.** 1 day
**C.** 7 days
**D.** 90 days
**Answer: (SHOW ANSWER)**
The Fail-safe retention period for transient and temporary tables in Snowflake is 0 days. Fail-safe is a feature designed to protect data against accidental loss or deletion by retaining historical data for a period after its Time Travel retention period expires. However, transient and temporary tables, which are designed for temporary or short-term storage and operations, do not have a Fail-safe period. Once the data is deleted or the table is dropped, it cannot be recovered.
References:
Snowflake Documentation: Understanding Fail-safe

**NEW QUESTION: 235**
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: 236**
How can a data provider validate that a secure view is configured to display only the data the provider wishes to expose?
**A.** Log in to the data consumer account and check if the secure view data is appearing as expected.
**B.** Create a data share for a test data consumer account and check if the secure view data is appearing as expected.
**C.** Query the secure view from a consumer account by setting the share_restrictions parameter.
**D.** Simulate querying the secure view by setting the simulated_data_sharing_consumer session parameter.
**Answer: (SHOW ANSWER)**
The most effective way for a data provider to validate secure view configurations is to create a data share for a test data consumer account. This method allows the provider to review and confirm that only the intended data is accessible in the secure view. Secure views are designed to mask or restrict data visibility, so creating a test share replicates the consumer's experience and ensures data security before sharing with actual consumers.

**NEW QUESTION: 237**
What are the benefits of using the Snowflake Access_History view? (Select TWO)
**A.** It allows the user to compare query performance data over time.
**B.** It identifies data that is not being used so that it can be archived or deleted.
**C.** It tracks how sensitive data has been moved.
**D.** It allows users to track the relationships among objects.
**E.** It audits whether query policies are being enforced correctly.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 238**
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: (SHOW ANSWER)**

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: 239**
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: 240**
When unloading data from Snowflake using the COPY INTO < location > command, what is a key consideration for handling NULL values?
Freecram.net
**A.** NULL values are automatically converted to empty strings.
**B.** A string value should be specified to represent NULL values using the NULL_IF option.
**C.** If a NULL value is detected during data loading, the process will stop and report an error.
**D.** NULL values are automatically skipped during the unload process.
**Answer: (SHOW ANSWER)**
The correct answer is B. A string value should be specified to represent NULL values using the NULL_IF option .

When unloading data from Snowflake into files using the COPY INTO < location > command, Snowflake needs a textual representation for SQL NULL values in the unloaded files. This can be controlled by the NULL_IF file format option.
Why B is correct:
The NULL_IF option specifies one or more string values that represent SQL NULL values. During unloading, Snowflake uses this option to determine how NULL values should appear in the output files.
For example, a file format could specify:
NULL_IF = ( ' NULL ' )
This means SQL NULL values can be written using the string NULL in the unloaded file.
Why the other options are incorrect:
A). Snowflake does not always automatically convert NULL values to empty strings in all unload scenarios.
The representation depends on file format settings.
C). NULL values do not cause the unload process to stop with an error by default.
D). NULL values are not automatically skipped during unloading. They are included in the output using the configured representation.
Official Snowflake documentation reference:
Snowflake documentation for file format options explains that NULL_IF specifies the string used to represent SQL NULL values when loading or unloading data.
Reference: Snowflake Documentation - COPY INTO < location > command; Snowflake Documentation - File format options, NULL_IF; SnowPro Core Study Guide - Data Loading and Unloading.
==

**NEW QUESTION: 241**
A JSON document is stored in the source_colum 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 aVARIANTcolumn, 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 thenamefrom the first element of theelementsarray, the correct syntax isSource_column:elements[0] .
name.
References: Snowflake Documentation on Semi-Structured Data

**Valid COF-C03 Dumps** shared by ExamDiscuss.com for Helping Passing COF-C03 Exam! ExamDiscuss.com now offer the **newest COF-C03 exam dumps**, the ExamDiscuss.com COF-C03 exam **questions have been updated**

and **answers have been corrected** get the **newest** ExamDiscuss.com COF-C03 dumps with Test Engine here: https://www.examdiscuss.com/Snowflake/exam/COF-C03/premium/ (**1187** Q&As Dumps, **35%OFF** **Special Discount Code:** **freecram**)

**NEW QUESTION: 242**
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: 243**
A user is setting up Snowpipe automation to load data from an Amazon S3 bucket into Snowflake. The user wants Snowflake to automatically start loading files as soon as new ones are added to the bucket. Which steps are required to configure this automation using cloud messaging? Select TWO.
**A.** Configure an S3 event notification to an Amazon Simple Notification Service, or SNS, topic.
**B.** Create a task that monitors the stage for new files.
**C.** Call the insertFiles REST API endpoint after each file upload.
**D.** Enable AUTO_INGEST on the Snowpipe definition.
**E.** Define an external table and use it as the source for a materialized view.
**Answer: (SHOW ANSWER)**
For Snowpipe auto-ingest on Amazon S3, cloud event notifications are configured so Snowflake is notified when new files arrive. The pipe must also be created with AUTO_INGEST = TRUE.
Why A and D are correct:
Amazon S3 event notifications can publish file arrival events through Amazon SNS/SQS integration, and Snowpipe uses those notifications when AUTO_INGEST is enabled.
Reference: Snowflake Documentation - Snowpipe auto-ingest on Amazon S3; CREATE PIPE.
==

**NEW QUESTION: 244**
What are the default Time Travel and Fail-safe retention periods for transient tables?
**A.** Time Travel - 1 day. Fail-safe - 1 day
**B.** Time Travel - 0 days. Fail-safe - 1 day
**C.** Time Travel - 1 day. Fail-safe - 0 days
**D.** Transient tables are retained in neither Fail-safe nor Time Travel
**Answer: C (LEAVE A REPLY)**

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

**NEW QUESTION: 245**
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: B,E (LEAVE A REPLY)**
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: 246**
When will Snowflake charge credits for the use of the Cloud Services layer?
**A.** Credits will be charged only when a virtual warehouse consumes serverless compute services.
**B.** Credits will be charged when the daily consumption of cloud services resources exceeds Wo of the daily warehouse usage.
**C.** Credits will be charged only when running a Snowflake-provisioned compute warehouse comfute_wh.
**D.** Credits will be charged whenever the Cloud Services layer is used.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 247**
What is the !define < variable > = < value > command used for in SnowSQL?
**A.** Assigns a name to a variable
**B.** Assigns a value to a variable
**C.** Sets the context for a variable
**D.** Establishes the metrics of a variable
**Answer: (SHOW ANSWER)**
The correct answer is B. Assigns a value to a variable .
Corrected wording:
The original command was written as idefine < variable > = < value > . In SnowSQL, the correct command is:
!define < variable > = < value >
SnowSQL supports variables that can be defined and reused in SQL scripts or interactive sessions. The !

define command assigns a value to a variable.
Why B is correct:
!define < variable > = < value > creates or sets a SnowSQL variable to the specified value.
Example:
!define table_name=my_table
SELECT *
FROM & table_name;
In this example, table_name is assigned the value my_table, and the variable can then be referenced later in the script.
Why the other options are incorrect:
A). The variable name is part of the command syntax, but the purpose of the command is to assign a value to that variable.
C). Context is usually set with commands such as USE DATABASE, USE SCHEMA, USE WAREHOUSE, or SnowSQL connection options.
D). The command does not establish metrics.
Official Snowflake documentation reference:
Snowflake documentation for SnowSQL explains that !define is used to define variables for use in SnowSQL sessions or scripts.
Reference: Snowflake Documentation - SnowSQL variables; Snowflake Documentation - SnowSQL commands; SnowPro Core Study Guide - Snowflake Account and Resource Management.

**NEW QUESTION: 248**
A Snowflake user is using the QUERY_HISTORY view in the ACCOUNT_USAGE schema to gather information about two queries that ran in the past 10 minutes, but the data is missing. What could cause this to occur?
**A.** The queries were run on a multi-cluster virtual warehouse.
**B.** The queries were able to take advantage of the results cache.
**C.** The ACCOUNT_USAGE schema does not provide real-time results.
**D.** The user does not have privileges to access the QUERY_HISTORY view.
**Answer: (SHOW ANSWER)**
The correct answer is C. The ACCOUNT_USAGE schema does not provide real-time results .
The ACCOUNT_USAGE schema contains historical metadata and usage information for the Snowflake account, but its views have latency. This means recently executed queries may not appear immediately.
Corrected wording:
The object is typically referred to as the QUERY_HISTORY view in the ACCOUNT_USAGE schema, not QUERY_HISTORY_VIEW.
Why C is correct:
If queries ran only 10 minutes ago, they may not yet appear in SNOWFLAKE.ACCOUNT_USAGE.
QUERY_HISTORY because ACCOUNT_USAGE views are not real-time. For more immediate query history, users can use the INFORMATION_SCHEMA.QUERY_HISTORY table function, subject to its scope and retention limits.
Why the other options are incorrect:
A). Running queries on a multi-cluster warehouse does not prevent query history from being recorded.
B). Queries using the result cache can still appear in query history.

D). If the user lacked privileges, they would generally be unable to query the view or would have limited visibility, but the key clue is that the queries ran in the past 10 minutes.
Official Snowflake documentation reference:
Snowflake documentation states that ACCOUNT_USAGE views can have latency and should not be expected to return real-time data.
Reference: Snowflake Documentation - ACCOUNT_USAGE schema; Snowflake Documentation - QUERY_HISTORY view; Snowflake Documentation - INFORMATION_SCHEMA.QUERY_HISTORY; SnowPro Core Study Guide - Account Usage and Monitoring.
==

**NEW QUESTION: 249**
What is the following SQL command used for?
Select * from table(validate(t1, job_id = > ' _last ' ));
**A.** To validate external table files in table t1 across all sessions
**B.** To validate task SQL statements against table t1 in the last 14 days
**C.** To validate a file for errors before it gets executed using a COPY command
**D.** To return errors from the last executed COPY command into table t1 in the current session
**Answer: D (LEAVE A REPLY)**
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 checksthe results of the most recent data load operation and provides details on any errors that occurred during that process1.

**NEW QUESTION: 250**
Which automatic Snowpipe feature ensures that the same data file is not loaded multiple times into a table?
Freecram.net
**A.** Snowpipe deletes files from the stage after they are loaded, so they cannot be loaded twice.
**B.** Snowpipe references the metadata of previously loaded files and will not load new files with the same metadata.
**C.** Snowpipe renames each file after loading, ensuring that the same file cannot be loaded more than once.
**D.** Snowpipe moves each file to a temporary stage after loading, so it cannot be loaded more than once.
**Answer: (SHOW ANSWER)**
The correct answer is B. Snowpipe references the metadata of previously loaded files and will not load new files with the same metadata .
Snowflake keeps load metadata for files loaded into a table. Snowpipe uses this metadata to prevent duplicate loading of the same file into the same table.
Why B is correct:
Snowpipe tracks metadata for loaded files, including information that identifies files already processed. If the same file is encountered again, Snowpipe can use the load metadata to avoid loading it again.
Why the other options are incorrect:
A). Snowpipe does not automatically delete files from the stage after loading.
C). Snowpipe does not automatically rename staged files after loading.
D). Snowpipe does not automatically move loaded files to a temporary stage.

Official Snowflake documentation reference:
Snowflake documentation explains that Snowpipe maintains load metadata and prevents the same files from being loaded multiple times into the same table by default.
Reference: Snowflake Documentation - Snowpipe; Snowflake Documentation - Load metadata; Snowflake Documentation - Data loading considerations; SnowPro Core Study Guide - Data Loading and Unloading.
==

**NEW QUESTION: 251**
A user wants to access files 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: (SHOW ANSWER)**
A Pre-signed URL should be used to access files stored in a Snowflake stage without requiring authentication into Snowflake. Pre-signed URLs are simple HTTPS URLs that provide temporary access to a file via a web browser, using a pre-signed access token. The expiration time for the access token is configurable, and this type of URL allows users or applications to directly access or download the fileswithout needing to authenticate into Snowflake5.
References: [COF-C03] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 252**
Which tasks will help optimize virtual warehouse costs? (Select TWO)
**A.** Enabling auto-suspend
**B.** Limiting statement queue times
**C.** Enabling multi-cluster warehouses
**D.** Limiting statement timeout seconds
**E.** Enabling auto-resume
**Answer: (SHOW ANSWER)**
* Auto-suspend ensures that a warehouse automatically suspends when idle, reducing compute costs.
* Statement timeout prevents long-running queries from consuming excessive credits.
Reference:
Snowflake Docs: Managing Warehouses
Snowflake Docs: Statement Timeouts

**NEW QUESTION: 253**
Which command will create an ARRAY output from inputs ' a ' and ' b ' ?
**A.** LISTAGG( ' a ' , ' b ' );
**B.** ARRAY_CONSTRUCT( ' a ' , ' b ' );
**C.** TO_ARRAY( ' a ' , ' b ' );
**D.** AS_ARRAY( ' a ' , ' b ' );

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 254**
Which columns are part of the result set of the Snowflake LATERAL FLATTEN command? (Choose two.)
**A.** CONTENT
**B.** PATH
**C.** BYTE_SIZE
**D.** INDEX
**E.** DATATYPE
**Answer: (SHOW ANSWER)**
The LATERAL FLATTEN command in Snowflake produces a result set that includes several columns, among which PATH and INDEX are includedPATH indicates the path to the element within a data structure that needs to be flattened, and INDEX represents the index of the element if it is an array2.

**NEW QUESTION: 255**
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.

**NEW QUESTION: 256**
Which Snowflake object can be accessed in he 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

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

https://www.examdiscuss.com/Snowflake/exam/COF-C03/premium/ (**1187** Q&As Dumps, **35%OFF** **Special Discount Code:** **freecram**)

**NEW QUESTION: 257**
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: 258**
How can a Snowflake user traverse semi-structured data?
**A.** Insert a colon (:) between the VARIANT column name and any first-level element.
**B.** Insert a colon (:) between the VARIANT column name and any second-level element. C. Insert a double colon (: :) between the VARIANT column name and any first-level element.
**C.** Insert a double colon (: :) between the VARIANT column name and any second-level element.
**Answer: (SHOW ANSWER)**
To traverse semi-structured data in Snowflake, a user can insert a colon (:) between the VARIANT column name and any first-level element. This path syntax is used to retrieve elements in a VARIANT column4.

**NEW QUESTION: 259**
What value provides information about disk usage for operations where intermediate results do not fit in memory in a Query Profile?
**A.** IO
**B.** Network
**C.** Pruning
**D.** Spilling
**Answer: (SHOW ANSWER)**
In Snowflake, when a query execution requires more memory than what is available, Snowflake handles these situations by spilling the intermediate results to disk. This process is known as "spilling." The Query Profile in Snowflake includes a metric that helps users identify when and how much data spilling occurs during the execution of a query. This information is crucial for optimizing queries as excessive spilling can significantly slow down query performance. The value that provides this information about disk usage due to intermediate results not fitting in memory is appropriately labeled as "Spilling" in the Query Profile.

References:
Snowflake Documentation on Query Profile and Performance: This section explains the various components of the query profile, including the spilling metric, which indicates disk usage for operations where intermediate results exceed available memory.

**NEW QUESTION: 260**
What privileges or user roles are required to get or request access to data that is available in the Snowflake Marketplace? (Select TWO).
Freecram.net
**A.** SECURITYADMIN
**B.** MANAGE SHARE
**C.** ACCOUNTADMIN
**D.** USERADMIN
**E.** IMPORT SHARE
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 261**
Which JSON paths are considered to be equivalent in Snowflake? (Choose two.)
**A.** src:customer.email
**B.** SRC:customer.email
**C.** src[ ' customer ' ][ ' EMAIL ' ]
**D.** SRC:Customer.Email
**E.** src[ ' CUSTOMER ' ][ ' Email ' ]
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 262**
Which Snowflake feature is used to track sensitive data to ensure compliance, identify data, and monitor resource usage?
**A.** Multi-Factor Authentication (MFA)
**B.** Data classification
**C.** Data quality monitoring
**D.** Object tagging
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 263**
Which Snowflake table supports unstructured data?
**A.** Directory
**B.** Transient
**C.** Temporary
**D.** Permanent
**Answer: (SHOW ANSWER)**

While Snowflake primarily deals with structured and semi-structured data, it also has the capability to handle unstructured data. Unstructured data can be stored in Snowflake using variants of SQL data types in tables, which can be permanent tables. These permanent tables, while traditionally used for structured or semi- structured data (like JSON, Avro, or Parquet), can also accommodate unstructured data in the form of binary formats or strings, offering flexibility in data storage and analysis. However, the management and querying of unstructured data in Snowflake may require additional considerations compared to structured data.
References: Snowflake Documentation on Data Types

**NEW QUESTION: 264**
What action will cause the cache associated with compute resources to be dropped?
**A.** Adding a new standard virtual warehouse that is the same size as a running 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.** Decreasing the size of a running standard virtual warehouse
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 265**
Which operator is used to evaluate query pruning efficiency by comparing partitions scanned to partitions total?
**A.** TableScan
**B.** IndexScan
**C.** Filter
**D.** ExternalScan
**Answer: A (LEAVE A REPLY)**

**NEW QUESTION: 266**
Which statements reflect valid commands when using secondary roles? (Select TWO).
**A.** Use SECONDARY ROLES RESUME
**B.** USE SECONDARY ROLES SUSPEND
**C.** USE SECONDARY RLES ALL
**D.** USE SECONDARY ROLES ADD < Role Name >
**E.** Use SECONDARY ROLES NONE
**Answer: (SHOW ANSWER)**
USE SECONDARY ROLES ALL: Activates all granted secondary roles for the current session.
USE SECONDARY ROLES NONE: Deactivates all active secondary roles for the current session.
Incorrect Commands: The options referencing " RESUME " , " SUSPEND " , and " ADD " are not valid commands in the context of secondary roles.
References:
Snowflake Documentation (USE SECONDARY ROLES): https://docs.snowflake.com/en/sql-reference/sql /use-secondary-roles.html

**NEW QUESTION: 267**

Which of the following commands cannot be used within a reader account?
**A.** CREATE SHARE
**B.** ALTER WAREHOUSE
**C.** DROP ROLE
**D.** SHOW SCHEMAS
**E.** DESCRBE TABLE
**Answer: A (LEAVE A REPLY)**
In Snowflake, a reader account is a type of account that is intended for consuming shared data rather than performing any data management or DDL operations. The CREATE SHARE command is used to share data from your account with another account, which is not a capability provided to reader accounts. Reader accounts are typically restricted from creating shares, as their primary purpose is to read shared data rather than to share it themselves.
References:
Snowflake Documentation on Reader Accounts
SnowPro Core Certification Study Guide

**NEW QUESTION: 268**
What are two ways to create and manage Data Shares in Snowflake? (Choose two.)
**A.** Via the Snowflake Web Interface (UI)
**B.** Via thedata_share=trueparameter
**C.** Via SQL commands
**D.** Via Virtual Warehouses
**Answer: (SHOW ANSWER)**
In Snowflake, Data Shares can be created and managed in two primary ways:
Via the Snowflake Web Interface (UI): Users can create and manage shares through the graphical interface provided by Snowflake, which allows for a user-friendly experience.
Via SQL commands: Snowflake also allows the creation and management of shares using SQL commands.
This method is more suited for users who prefer scripting or need to automate the process.
Reference:https://docs.snowflake.com/en/user-guide/data-sharing-provider.html

**NEW QUESTION: 269**
Which Snowflake feature records changes mace to a table so actions can be taken using that change data capture?
**A.** Materialized View
**B.** Pipe
**C.** Stream
**D.** Task
**Answer: (SHOW ANSWER)**
Snowflake ' s Streams feature is specifically designed for change data capture (CDC). A stream records insert, update, and delete operations performed on a table, and allows users to query these changes. This enables actions to be taken on the changed data, facilitating processes like incremental data loads and real-time analytics. Streams provide a powerful mechanism for applications to respond to data changes in Snowflake tables efficiently.
References: Snowflake Documentation on Streams

**NEW QUESTION: 270**

A query is using more credits than expected. The Query Profile shows that a majority of the query execution is spent on remote disk I/O.
How can this be prevented m the future?

Freecram.net

**A.** Increase the number of clusters in the virtual warehouse.
**B.** Increase the size of the virtual warehouse.
**C.** Convert the virtual warehouse to Maximized mode.
**D.** Increase the auto suspend time in the virtual warehouse.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 271**

Which function determines the kind of value stored in a VARIANT column?
**A.** CHECK_JSON
**B.** IS_ARRAY
**C.** IS_JSON
**D.** TYPEOF
**Answer: (SHOW ANSWER)**

The function used to determine the kind of value stored in a VARIANT column in Snowflake is TYPEOF.
Understanding VARIANT Data Type:
VARIANT is a flexible data type in Snowflake that can store semi-structured data, such as JSON, Avro, and XML.
This data type can hold values of different types, including strings, numbers, objects, arrays, and more.
Using TYPEOF Function:
The TYPEOF function returns the data type of the value stored in a VARIANT column.
It helps in identifying the type of data, which is crucial for processing and transforming semi-structured data accurately.
Example Usage:
SELECT TYPEOF(variant_column)
FROM my_table;
This query retrieves the type of data stored invariant_columnfor each row inmy_table.
Possible return values include ' OBJECT ' , ' ARRAY ' , ' STRING ' , ' NUMBER ' , etc.
Benefits:
Simplifies data processing: Knowing the data type helps in applying appropriate transformations and validations.
Enhances query accuracy: Ensures that operations on VARIANT columns are performed correctly based on the data type.
Snowflake Documentation: TYPEOF
Snowflake Documentation: VARIANT Data Type

**Valid COF-C03 Dumps** shared by ExamDiscuss.com for Helping Passing COF-C03 Exam! ExamDiscuss.com now offer the **newest COF-C03 exam dumps**, the ExamDiscuss.com COF-C03 exam **questions have been updated**

and **answers have been corrected** get the **newest** ExamDiscuss.com COF-C03 dumps with Test Engine here: https://www.examdiscuss.com/Snowflake/exam/COF-C03/premium/ (**1187** Q&As Dumps, **35%OFF** **Special Discount Code:** **freecram**)

**NEW QUESTION: 272**
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: 273**
How does Snowflake support the use of structured and semi-structured data? (Select TWO)
**A.** Structured data is supported through the use of traditional relational database tables, organized into rows and columns.
**B.** Structured and semi-structured data are stored in separate physical storage locations; users must choose which type of data to use when creating a table.
**C.** Semi-structured data is supported through the use of automatically indexed database tables.
**D.** All semi-structured data is automatically converted into structured data when it is loaded into Snowflake.
**E.** Users can load and query semi-structured data using SQL using the SQL semi-structured data handling functions.
**Answer: (SHOW ANSWER)**
Snowflake supports structured data in standard tables and semi-structured data in VARIANT, OBJECT, and ARRAY columns.
Users can load and query both using standard SQL functions , without schema constraints for semi- structured formats.
Reference:
Snowflake Docs: Semi-Structured Data Overview

**NEW QUESTION: 274**
What are valid values for the FIELD_OPTIONALLY_ENCLOSED_BYoption in the copy into < location > command used during data unloading? (Select TWO).
**A.** ' NULL '
**B.** NONE
**C.** Single quote character ( ' )

**D.** NULL
**E.** ' NONE '
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 275**
Which of the following objects are contained within a schema? (Choose two.)
**A.** Role
**B.** Stream
**C.** Warehouse
**D.** External table
**E.** User
**F.** Share
**Answer: (SHOW ANSWER)**
In Snowflake, a schema is a logical grouping of database objects, which can include streams and external tables. A stream is an object that allows users to query data that has changed in specified tables or views, and an external table is a table that references data stored outside of Snowflake. Roles, warehouses, users, and shares are not contained within a schema. References: SHOW OBJECTS, Database, Schema, & Share DDL

**NEW QUESTION: 276**
A Snowflake user has two tables that contain numeric values and is trying to find out which values are present in both tables. Which set operator should be used?
**A.** INTERSECT
**B.** MFRCK
**C.** MINUS
**D.** UNION
**Answer: (SHOW ANSWER)**
To find out which numeric values are present in both tables, the INTERSECT set operator should be used.This operator returns rows from one query's result set which also appear in another query's result set, effectively finding the common elements between the two tables45.

**NEW QUESTION: 277**
What MINIMUM permissions are required to create a pipe in Snowflake? (Select TWO).
**A.** SELECT and INSERT on the table in the pipe definition
**B.** CREATE PIPE at the schema level
**C.** CREATE SCHEMA at the database level
**D.** OWNERSHIP on the table in the pipe definition
**E.** USAGE on the virtual warehouse
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 278**
Which Snowflake objects are supported by Snowpipe Streaming? (Select TWO)

**A.** Tables using sequences to generate names for columns
**B.** Tables with columns containing the GEOMETRY datatype
**C.** Temporary tables
**D.** Permanent tables
**E.** Tables with a default column that is not NULL
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 279**
Which formats does Snowflake store unstructured data in? (Choose two.)
**A.** GeoJSON
**B.** Array
**C.** XML
**D.** Object
**E.** BLOB
**Answer: A,C (LEAVE A REPLY)**
Snowflake supports storing unstructured data and provides native support for semi-structured file formats such as JSON, Avro, Parquet, ORC, and XML1. GeoJSON, being a type of JSON, and XML are among the formats that can be stored in Snowflake. References: [COF-C03] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 280**
A developer expects Snowflake to refresh a dynamic table incrementally, but Snowflake selects full refresh mode instead.
Which query operator can prevent incremental refresh and cause Snowflake to use full refresh mode when REFRESH_MODE = AUTO?
Freecram.net
**A.** DISTINCT
**B.** GROUP BY
**C.** LATERAL JOIN
**D.** UNION ALL
**Answer: C (LEAVE A REPLY)**
Snowflake supports incremental refresh for many common query constructs, including supported uses of DISTINCT, GROUP BY, and UNION ALL.
General lateral joins are not supported for incremental dynamic-table refresh. A documented exception exists for supported uses of LATERAL FLATTEN.
When a dynamic table is created with REFRESH_MODE = AUTO, Snowflake analyzes the defining query.
When the query cannot be refreshed incrementally, Snowflake can choose full refresh mode.
When REFRESH_MODE = INCREMENTAL is explicitly requested and the query contains an unsupported construct, dynamic-table creation can fail instead of silently changing the requested refresh mode.
References:
Snowflake Documentation - Supported queries for incremental refresh: https://docs.snowflake.com/en/user-guide/dynamic-tables-supported-queries Snowflake Documentation - Dynamic-table refresh modes: https://docs.snowflake.com/en/user-guide

/dynamic-tables-refresh
==

**NEW QUESTION: 281**
What data-type constructs are used to create a hierarchy when organizing semi-structured data? (Select TWO).
**A.** An ARRAY or OBJECT holds a value of type INTEGER.
**B.** A VARIANT can hold a value of any other data type.
**C.** An ARRAY or OBJECT holds a value of type VECTOR.
**D.** An ARRAY or OBJECT holds a value of type VARIANT.
**E.** An ARRAY or OBJECT holds a value of type VARCHAR.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 282**
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: 283**
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: 284**

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: 285**
When should a stored procedure be created with caller ' s rights?
**A.** When the caller needs to be prevented from viewing the source code of the stored procedure
**B.** When the caller needs to run a statement that could not execute outside of the stored procedure
**C.** When the stored procedure needs to run with the privileges of the role that called the stored procedure
**D.** When the stored procedure needs to operate on objects that the caller does not have privileges on
**Answer: (SHOW ANSWER)**
Stored procedures in Snowflake can be created with either ' owner ' s rights ' or ' caller ' s rights ' . A stored procedure created with caller ' s rights executes with the privileges of the role that calls the procedure, not the privileges of the role that owns the procedure. This is particularly useful in scenarios where the procedure needs to perform operations that depend on the caller ' s access permissions, ensuring that the procedure can only access objects that the caller is authorized to access.
Reference: Snowflake ' s official documentation on stored procedures, specifically the section on execution context (caller ' s rights vs. owner ' s rights), provides detailed guidance on when to use caller ' s rights: https://docs.snowflake.com/en/sql-reference/stored-procedures-usage.html#caller-s-rights-vs-owner-s-rights

**NEW QUESTION: 286**
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.
==

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

**NEW QUESTION: 287**
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: 288**
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: A (LEAVE A REPLY)**
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: 289**
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: 290**
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: 291**
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: 292**

At what levels can a network policy be configured in Snowflake? (Select TWO).
**A.** User level
**B.** Virtual warehouse level
**C.** Database level
**D.** Account level
**E.** Role level
**Answer: A,D (LEAVE A REPLY)**

**NEW QUESTION: 293**
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 returnvalues and cannot execute database operations1.

**NEW QUESTION: 294**
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: 295**
Which steps will help optimize query performance? (Select TWO).
**A.** Indexing a column
**B.** Using the query acceleration service

**C.** Decreasing the size of the virtual warehouse
**D.** Clustering a table
**E.** Increasing the size of the micro-partitions
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 296**
Which file function gives a user or application access to download unstructured data from a Snowflake stage?
**A.** BUILD_SCOPED_FILE_URL
**B.** BUILD_STAGE_FILE_URL
**C.** GET_PRESIGNED_URL
**D.** GET STAGE LOCATION
**Answer: (SHOW ANSWER)**
The function that provides access to download unstructured data from a Snowflake stage is:
C). GET_PRESIGNED_URL: This function generates a presigned URL for a single file within a stage. The generated URL can be used to directly access or download the file without needing to go through Snowflake.
This is particularly useful for unstructured data such as images, videos, or large text files, where direct access via a URL is needed outside of the Snowflake environment.
Example usage:
SELECTGET_PRESIGNED_URL( ' stage_name ' , ' file_path ' );
This function simplifies the process of securely sharing or accessing files stored in Snowflake stages with external systems or users.
References:
Snowflake Documentation: GET_PRESIGNED_URL Function at Snowflake Documentation

**NEW QUESTION: 297**
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: 298**
Which Snowflake objects can be secured to protect data privacy? (Select TWO).
**A.** Materialized views
**B.** Tags
**C.** Schemas

**D.** Stages
**E.** Views
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 299**
Which resource monitor setting will cancel all active queries in a virtual warehouse when the threshold is met?
**A.** SUSPEND
**B.** SUSPEND_IMMEDIATE
**C.** NOTIF
**D.** NOTIFY_USERS
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 300**
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: 301**
Which table function is used to perform additional processing on the results of a previously-run query?
**A.** QUERY_HISTORY
**B.** RESULT_SCAN
**C.** DESCRIBE_RESULTS
**D.** QUERY HISTORY BY SESSION
**Answer: (SHOW ANSWER)**
TheRESULT_SCANtable function is used in Snowflake to perform additional processing on the results of a previously-run query. It allows users to reference the result set of a previous query by its query ID, enabling further analysis or transformations without re-executing the original query.
Snowflake Documentation: RESULT_SCAN

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

**NEW QUESTION: 302**
When cloning a schema, which Snowflake object will not be included in the clone?
**A.** A named internal stage
**B.** An external stage
**C.** A task
**D.** A User-Defined Function (UDF)
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 303**
How does Snowflake Copilot interact with customer data?
**A.** Copilot analyzes customer data to optimize query performance.
**B.** Copilot provides assistance when creating and interpreting queries.
**C.** Copilot rewrites table schemas based on inferred patterns in customer datasets.
**D.** Copilot executes stored procedures in response to natural language prompts.
**Answer: (SHOW ANSWER)**
The correct answer is B. Copilot provides assistance when creating and interpreting queries .
Snowflake Copilot is an AI-assisted feature that helps users work with Snowflake more efficiently. It can assist with writing SQL, understanding queries, and interpreting query-related information.
Why B is correct:
Copilot is designed to help users create, refine, and understand SQL queries using natural language assistance.
Why the other options are incorrect:
A). Query optimization is primarily handled by Snowflake's optimizer and performance features, not Copilot.
C). Copilot does not automatically rewrite table schemas based on inferred data patterns.
D). Copilot is not primarily a feature for executing stored procedures from prompts.
Official Snowflake documentation reference:
Snowflake documentation describes Snowflake Copilot as an AI assistant for helping users generate, understand, and work with SQL and Snowflake data workflows.
Reference: Snowflake Documentation - Snowflake Copilot; Snowflake Cortex; SnowPro Core Study Guide - Snowflake Architecture.
==

**NEW QUESTION: 304**
Which database objects can be shared with the Snowflake secure data sharing feature? (Choose two.)
**A.** Files

**B.** External tables
**C.** Secure User-Defined Functions (UDFs)
**D.** Sequences
**E.** Streams
**Answer: (SHOW ANSWER)**
Snowflake's secure data sharing feature allows sharing of certain database objects with other Snowflake accounts.Among the options provided, external tables and secure UDFs can be shared

**NEW QUESTION: 305**
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: 306**
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.
==

**NEW QUESTION: 307**
Which Snowflake feature or tool helps troubleshoot issues in SQL query expressions that commonly cause performance bottlenecks?
**A.** Persisted query results
**B.** QUERY_HISTORY View
**C.** Query acceleration service
**D.** Query Profile
**Answer: (SHOW ANSWER)**
The Snowflake feature that helps troubleshoot issues in SQL query expressions and commonly identify performance bottlenecks is theQuery Profile. The Query Profile provides a detailed breakdown of a query ' s execution plan, including each operation ' s time and resources consumed. It visualizes the steps involved in the query execution, highlighting areas that may be causing inefficiencies, such as full table scans, large joins, or operations that could benefit from optimization.

By examining the Query Profile, developers and database administrators can identify and troubleshoot performance issues, optimize query structures, and make informed decisions about potential schema or indexing changes to improve performance.
References:
Snowflake Documentation on Query Profile: Using the Query Profile

**NEW QUESTION: 308**
While clustering a table, columns with which data types can be used as clustering keys? (Select TWO).
**A.** BINARY
**B.** GEOGRAPHY
**C.** GEOMETRY
**D.** OBJECT
**E.** VARIANT
**Answer: (SHOW ANSWER)**
A clustering key can be defined when a table is created by appending a CLUSTER Where each clustering key consists of one or more table columns/expressions, which can be of any data type, except GEOGRAPHY, VARIANT, OBJECT, or ARRAY https://docs.snowflake.com/en/user-guide/tables-clustering-keys

**NEW QUESTION: 309**
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: 310**
Which type of charts are supported by Snowsight? (Select TWO)
**A.** Pie charts
**B.** Line charts
**C.** Gantt charts
**D.** Scatterplots
**E.** Flowcharts
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 311**
A user has an application that writes a new Tile to a cloud storage location every 5 minutes.
What would be the MOST efficient way to get the files into Snowflake?
**A.** Create a task that runs a copy into operation from an external stage every 5 minutes
**B.** Create a task that puts the files in an internal stage and automate the data loading wizard
**C.** Create a task that runs a GET operation to intermittently check for new files
**D.** Set up cloud provider notifications on the Tile location and use Snowpipe with auto-ingest
**Answer: (SHOW ANSWER)**
The most efficient way to get files into Snowflake, especially when new files are being written to a cloud storage location at frequent intervals, is to use Snowpipe with auto-ingest. Snowpipe is Snowflake's continuous data ingestion service that loads data as soon as it becomes available in a cloud storage location.
By setting up cloud provider notifications, Snowpipe can be triggered automatically whenever new files are written to the storage location, ensuring that the data is loaded into Snowflake with minimal latency and without the need for manual intervention or scheduling frequent tasks.
References:
Snowflake Documentation on Snowpipe
SnowPro Core Certification Study Guide

**NEW QUESTION: 312**
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: 313**
A Snowflake user is investigating a slow-running query using the Query Profile. What would be a reason to increase the size of the virtual warehouse?
**A.** Percentage Scanned from Cache: 81.22%
**B.** Processing: 51.8%
**C.** Remote Disk I/O: 40.6%

**D.** Bytes Scanned: 3.24 GB

**Answer: B (LEAVE A REPLY)**

The correct answer is B. Processing: 51.8% .

A high percentage of query time spent in Processing indicates that the query is spending significant time on CPU-intensive work, such as joins, aggregations, sorting, window functions, or expression evaluation.

Increasing the virtual warehouse size can provide more compute resources and may reduce processing time.

Why B is correct:

If the Query Profile shows a large portion of time spent in processing, scaling up the warehouse can help because a larger warehouse provides more compute power.

Why the other options are incorrect:

A). A high percentage scanned from cache is generally positive. It means much of the data was served from cache instead of remote storage.

C). High remote disk I/O may indicate data is being read from remote storage. Increasing warehouse size may not be the first or best fix; improving pruning, clustering, filters, or reducing scanned data may be more appropriate.

D). Bytes Scanned: 3.24 GB by itself is not enough to justify increasing warehouse size. The reason to scale up is usually compute pressure, not simply a raw scan number.

Official Snowflake documentation reference:

Snowflake documentation explains that Query Profile helps identify time spent in processing, local disk I/O, remote disk I/O, synchronization, and other execution categories. Increasing warehouse size can help when queries are constrained by compute processing.

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

**NEW QUESTION: 314**

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: 315**

Which database object can be created as temporary for a given session?

**A.** Stored procedure

**B.** Stage

**C.** User-Defined Function (UDF)

**D.** View

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 316**
Which Snowflow object does not consume and storage costs?
**A.** Secure view
**B.** Materialized view
**C.** Temporary table
**D.** Transient table
**Answer: (SHOW ANSWER)**
Temporary tables in Snowflake do not consume storage costs. They are designed for transient data that is needed only for the duration of a session. Data stored in temporary tables is held in the virtual warehouse's cache and does not persist beyond the session's lifetime, thereby not incurring any storage charges.
References:
Snowflake Documentation: Temporary Tables

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

**NEW QUESTION: 317**
When unloading data, which file format will keep floating-point number column data without truncating it?
**A.** Parquet
**B.** XML
**C.** JSON
**D.** CSV
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 318**
Which function will return the URL for an external stage by using the stage name as the input?
**A.** GET_ABSOLUTE_PATH
**B.** GET_PRESIGNED_URL
**C.** GET_STAGE_LOCATION
**D.** BUILD_STAGE_FILE_URL
**Answer: (SHOW ANSWER)**
The GET_STAGE_LOCATION function retrieves the full external URL associated with a named stage. It is useful for accessing the storage location of external stages configured in Snowflake.
Reference:

Snowflake Docs: GET_STAGE_LOCATION

**NEW QUESTION: 319**
Snowflake supports the use of external stages with which cloud platforms? (Choose three.)
**A.** Amazon Web Services
**B.** Docker
**C.** IBM Cloud
**D.** Microsoft Azure Cloud
**E.** Google Cloud Platform
**F.** Oracle Cloud
**Answer: (SHOW ANSWER)**
Snowflake supports the use of external stages with Amazon Web Services (AWS), Microsoft Azure Cloud, and Google Cloud Platform (GCP).These platforms allow users to stage data externally and integrate with Snowflake for data loading operations

**NEW QUESTION: 320**
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?
Freecram.net
**A.** GET
**B.** PUT
**C.** CREATE STREAM
**D.** COPY INTO < location >
**Answer: B (LEAVE A REPLY)**
To upload data files from a local directory or folder on a client machine to an internal stage in Snowflake, thePUTcommand is used. ThePUTcommand 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 theCOPY INTOcommand.
Reference: Snowflake Documentation on Using the PUT Command: https://docs.snowflake.com/en/sql-reference/sql/put.html

**NEW QUESTION: 321**
Which command can be used to unload data into an external named stage in Snowflake?
**A.** PUT
**B.** CREATE <object>
**C.** COPY INTO <table>
**D.** COPY INTO <location>
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 322**
Which of the following describes how clustering keys work in Snowflake?
**A.** Clustering keys update the micro-partitions in place with a full sort, and impact the DML operations.
**B.** Clustering keys sort the designated columns over time, without blocking DML operations
**C.** Clustering keys create a distributed, parallel data structure of pointers to a table ' s rows and columns
**D.** Clustering keys establish a hashed key on each node of a virtual warehouse to optimize joins at run-time
**Answer: (SHOW ANSWER)**
Clustering keys in Snowflake work by sorting the designated columns over time. This process is done in the background and does not block data manipulation language (DML) operations, allowing for normal database operations to continue without interruption. The purpose of clustering keys is to organize the data within micro-partitions to optimizequery performance1.
References:
[ COF-C03 ] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Clustering1

**NEW QUESTION: 323**
What happens to the privileges granted to Snowflake system-defined roles?
**A.** The privileges cannot be revoked.
**B.** The privileges can be revoked by an ACCOUNTADMIN.
**C.** The privileges can be revoked by an orgadmin.
**D.** The privileges can be revoked by any user-defined role with appropriate privileges.
**Answer: A (LEAVE A REPLY)**
The privileges granted to Snowflake's system-defined roles cannot be revoked. System-defined roles, such as SYSADMIN, ACCOUNTADMIN, SECURITYADMIN, and others, come with a set of predefined privileges that are essential for the roles to function correctly within the Snowflake environment. These privileges are intrinsic to the roles and ensure that users assigned these roles can perform the necessary tasks and operations relevant to their responsibilities.
The design of Snowflake's role-based access control (RBAC) model ensures that system-defined roles have a specific set of non-revocable privileges to maintain the security, integrity, and operational efficiency of the Snowflake environment. This approach prevents accidental or intentional modification of privileges that could disrupt the essential functions or compromise the security of the Snowflake account.
References:
Snowflake Documentation on Access Control: Understanding Role-Based Access Control (RBAC)

**NEW QUESTION: 324**
A user needs to create a materialized view in the schema MYDB.MYSCHEMA. Which statements will provide this access?
**A.** GRANT ROLE MYROLE TO USER USER1;GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO ROLE MYROLE;
**B.** GRANT ROLE MYROLE TO USER USER1;GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER USER1;

**C.** GRANT ROLE MYROLE TO USER USER1;GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB. K " - ' SCHEMA TO USER! ;
**D.** GRANT ROLE MYROLE TO USER USER1;GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO MYROLE;
**Answer: (SHOW ANSWER)**
To provide a user with the necessary access to create a materialized view in a schema, the user must be granted a role that has the CREATE MATERIALIZED VIEW privilege on that schema. First, the role is granted to the user, and then the privilege is granted to the role

**NEW QUESTION: 325**
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)**
TheSearch Optimization Servicein Snowflake enhances query performance by creating apersistent data structurethat enables faster access to specific data, particularly for queries with selective filters oncolumns 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: 326**
Which URL type allows users to access unstructured data without authenticating into Snowflake or passing an authorization token?
**A.** Pre-signed URL
**B.** Scoped URL
**C.** Signed URL
**D.** File URL
**Answer: (SHOW ANSWER)**
Pre-signed URLs in Snowflake allow users to access unstructured data without the need for authentication into Snowflake or passing an authorization token.These URLs are open and can be directly accessed or downloaded by any user or application, making them ideal for business intelligence applications or reporting tools that need to display unstructured file contents

**NEW QUESTION: 327**
A company configures two resource monitors:
* An account-level resource monitor with a quota of 1,000 credits and the action NOTIFY & SUSPEND.
* A warehouse-level resource monitor for the Marketing warehouse with a quota of 200 credits and the action NOTIFY & SUSPEND.

What happens when the Marketing warehouse reaches 200 credits while total account usage is only 700 credits?

**A.** The Marketing warehouse is suspended immediately, and all running queries are canceled.

**B.** A notification is sent, and the Marketing warehouse is suspended after currently running statements finish.

**C.** A notification is sent, and the warehouse is suspended only when no ACCOUNTADMIN queries are running.

**D.** The Marketing warehouse continues running until the account-level resource monitor reaches 1,000 credits.

**Answer: (SHOW ANSWER)**

A warehouse can be governed by its assigned warehouse-level resource monitor independently of an account- level resource monitor.

When the Marketing warehouse reaches its own 200-credit threshold, the NOTIFY & SUSPEND action takes effect even though the account-level monitor has not reached its 1,000-credit threshold.

The suspend action prevents new statements from starting and suspends the warehouse after currently running statements complete. It does not immediately cancel running queries.

Immediate cancellation corresponds to the NOTIFY & SUSPEND IMMEDIATE action. The user's administrative role does not exempt a running query from the monitor's configured threshold action.

References:

Snowflake Documentation - Resource monitors: https://docs.snowflake.com/en/user-guide/resource-monitors

Snowflake SQL Reference - CREATE RESOURCE MONITOR: https://docs.snowflake.com/en/sql- reference/sql/create-resource-monitor

**NEW QUESTION: 328**

Assuming the data provider has granted the required permissions to the objects in a share, which statement will grant permissions to the role consumer_role, allowing access to the database share_db?

**A.** grant all on database share_db to role consumer_role;

**B.** grant ownership on database share_db to role consumer_role;

**C.** grant imported privileges on database share_db to role consumer_role;

**D.** grant usage on database share_db to role consumer_role;

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 329**

Which URL type permits temporary access to a staged file without granting privileges to the stage?

**A.** Build_Stage_File_URL

**B.** File URL

**C.** Scoped URL

**D.** Pre-signed URL

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 330**

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?

Freecram.net

**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: 331**
Which Snowflake database object can be used to track data changes made to table data?
**A.** Tag
**B.** Task
**C.** Stream
**D.** Stored procedure
**Answer: (SHOW ANSWER)**
A Stream object in Snowflake is used for change data capture (CDC), which records data manipulation language (DML) changes made to tables, including inserts, updates, and deletes3.

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

**NEW QUESTION: 332**
Which feature allows a user to review the monthly projected costs of a Snowflake resource against the monthly actual costs?
**A.** Budgets
**B.** Auto-suspend
**C.** Scaling policy
**D.** Resource monitor
**Answer: (SHOW ANSWER)**
Snowflake Budgets is the cost-management feature designed around a monthly spending limit. It tracks credit consumption for an account or a defined group of supported objects, presents spending against the configured monthly budget, and forecasts whether current usage is likely to exceed that limit.
Notifications and custom actions can be based on projected or actual spending. Auto-suspend controls warehouse idle time, a scaling policy governs multi-cluster behavior, and a resource monitor tracks warehouse credit usage and can

notify or suspend warehouses at thresholds. Those features do not provide the same monthly actual-versus-projected budget view. Therefore, Budgets is the appropriate feature for reviewing current spending and projected monthly cost for a Snowflake resource or resource group.

**NEW QUESTION: 333**
How can a row access policy be applied to a table or a view? (Choose two.)
**A.** Within the policy DDL
**B.** Within the create table or create view DDL
**C.** By future APPLY for all objects in a schema
**D.** Within a control table
**E.** Using the command ALTER < object > ADD ROW ACCESS POLICY < policy > ;
**Answer: A,E (LEAVE A REPLY)**
A row access policy can be applied to a table or a view within the policy DDL when defining the policy. Additionally, an existing row access policy can be applied to a table or a view using the ALTER < object > ADD ROW ACCESS POLICY < policy > command

**NEW QUESTION: 334**
Which feature or service should be used to identify data about a department that was responsible for a recent revenue shortfall?
**A.** Snowpark
**B.** Snowpipe
**C.** ML Functions
**D.** AI_EXTRACT
**Answer: C (LEAVE A REPLY)**
Snowflake ML Functions include Top Insights, which performs key-driver analysis to explain why a business metric changed. It can compare groups or time periods and rank dimensions that contributed most strongly to a result. Snowflake documentation uses the example of identifying locations, salespeople, customers, verticals, and other factors responsible for a recent revenue shortfall. A department can be evaluated as one of those dimensions, allowing the analysis to identify which department contributed to the decline. Snowpark is a developer framework for processing data with languages such as Python, Java, and Scala. Snowpipe loads files continuously, while AI_EXTRACT extracts structured information from documents. Those features do not directly perform the requested driver analysis. Therefore, ML Functions is the correct answer.

**NEW QUESTION: 335**
What operations can be performed while loading a simple CSV file into a Snowflake table using the COPY INTO command? (Select TWO).
**A.** Performing aggregate calculations
**B.** Reordering the columns
**C.** Grouping by operations
**D.** Converting the datatypes
**E.** Selecting the first few rows

**Answer: (SHOW ANSWER)**

When loading a simple CSV file into a Snowflake table using the COPY INTO command, you can perform various transformations and adjustments on the data as part of the loading process. Specifically, you can:

B). Reorder the columns: Specify the order of columns in the COPY INTO command to match the structure of the target table if it differs from the order of columns in the source CSV file.

D). Convert the datatypes: Explicitly convert the datatypes of the data being loaded to match the datatypes of the columns in the target table. This can be necessary when the source data ' s format does not match the target table ' s expected datatype.

References:

Snowflake Documentation: Using the COPY INTO Command for Data Loading

**NEW QUESTION: 336**

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: C (LEAVE A REPLY)**

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: 337**
A size Small standard virtual warehouse running in auto-suspend mode is used for a production Business Intelligence (BI) reporting environment. The BI reports generate quickly in the early morning, but as the day progresses and there are more users, it takes longer to run the reports. How can this situation be addressed?
**A.** Disable the auto-suspend mode.
**B.** Change the standard warehouse to a multi-cluster warehouse.
**C.** Increase the size of the virtual warehouse.
**D.** Enable the query acceleration service.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 338**
Which type of join will list a I rows in the specified table, even if those rows have no match in the other table?
**A.** Cross join
**B.** Inner join
**C.** Natural join
**D.** Outer join
**Answer: (SHOW ANSWER)**
An outer join, specifically a left outer join, will list all rows from the left table and match them with rows from the right table. If there is no match, the result will still include the row from the left table, with NULLs for columns from the right table.References:Based on general SQL knowledge as of 2021.

**NEW QUESTION: 339**
How can a Data Exchange Administrator provide a user with account access to a Data Exchange?
**A.** Add the user to the Data Exchange.
**B.** Grant the user the USERADMIN role.
**C.** Enable the IMPORT SHARE privilege and grant this privilege to the user.
**D.** Create a new database for the Data Exchange and provide access to the user.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 340**
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: (SHOW ANSWER)**
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: 341**
When retrieving data from a table which type of query will benefit the MOST from use of the search optimization service?
**A.** Queries that involve one or more columns inside the cluster key using the lowest cardinality to select filters.
**B.** Queries on fields in a VARIANT column that will return all data as structured data.
**C.** Queries with highly-selective table filters that return a small subset of rows.
**D.** Queries with unpredictable data volumes.
**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 342**
Which column in the ACCESS_HISTORY view in Snowflake describes access history for a column and provides additional information on how columns have been altered over time?
**A.** object_modified_by_dd1
**B.** direct_objects_accessed
**C.** objects_modified
**D.** base_objects_accessed
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 343**
What are common issues found by using the Query Profile? (Choose two.)
**A.** Identifying queries that will likely run very slowly before executing them
**B.** Locating queries that consume a high amount of credits
**C.** Identifying logical issues with the queries
**D.** Identifying inefficient micro-partition pruning
**E.** Data spilling to a local or remote disk
**Answer: (SHOW ANSWER)**
The Query Profile in Snowflake is used to identify performance issues with queries. Common issues that can be found using the Query Profile include identifying inefficient micro-partition pruning (D) and data spilling to a local or remote disk (E).Micro-partition pruning is related to the efficiency of query execution, and data spilling occurs when the memory is insufficient, causing the query to write data to disk, which can slow down the query performance1.

**NEW QUESTION: 344**
Which of the following SQL statements will list the version of the drivers currently being used?
**A.** Execute SELECT CURRENT_ODBC_CLIENT(); from the Web UI
**B.** Execute SELECT CURRENT_JDBC_VERSION() ; from SnowSQL

**C.** Execute SELECT CURRENT_CLIENT(); from an application
**D.** Execute SELECT CURRENT_VERSION (); from the Python Connector
**Answer: (SHOW ANSWER)**
The correct SQL statement to list the version of the JDBC (Java Database Connectivity) drivers currently being used is to executeSELECT CURRENT_JDBC_VERSION();from within SnowSQL or any client application that utilizes JDBC for connecting to Snowflake. Snowflake provides specific functions to query the version of the client drivers or connectors being used, such as JDBC, ODBC, and others.
TheCURRENT_JDBC_VERSION()function is designed specifically to return the version of the JDBC driver in use.
It ' s important to note that Snowflake supports various types of drivers and connectors for connecting to different client applications, including ODBC, JDBC, Python, and others. Each of these connectors has its own method or function for querying the current version in use. For JDBC, the appropriate function isCURRENT_JDBC_VERSION(), reflecting the specificity required to obtain version information relevant to the JDBC driver specifically.
References:
Snowflake Documentation on Client Functions: This information can typically be found in the Snowflake documentation under the section that covers SQL functions, particularly those functions that provide information about the client or session.

**NEW QUESTION: 345**
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: 346**
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 executeCREATE < 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

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

**NEW QUESTION: 347**
Which system_defined Snowflake role has permission to rename an account and specify whether the original URL can be used to access the renamed account?
**A.** ACCOUNTADMIN
**B.** SECURITYADMIN
**C.** SYSADMIN
**D.** ORGADMIN
**Answer: (SHOW ANSWER)**
TheACCOUNTADMINrole in Snowflake has the highest level of privileges, including the ability to manage accounts, users, roles, and all objects within the account. This role is specifically granted the permission to rename an account and specify whether the original URL can be used to access the renamed account.
TheACCOUNTADMINrole encompasses broad administrative capabilities, ensuring that users assigned this role can perform critical account management tasks.
References: Snowflake Documentation on Roles and Permissions

**NEW QUESTION: 348**
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: 349**

Why would the SINGLE = TRUE option be included in a COPY INTO < location > statement?

**A.** To set the file format to one data type

**B.** To unload data to one output file

**C.** To limit a user ' s ability to split data into multiple files

**D.** To ensure that there is one attempt to load the data

**Answer: (SHOW ANSWER)**

The correct answer is B. To unload data to one output file .

The COPY INTO < location > command unloads data from a Snowflake table or query result into files in a stage. By default, Snowflake may unload data into multiple files depending on the amount of data and warehouse parallelism.

Why B is correct:

The SINGLE = TRUE option tells Snowflake to unload the result into a single file instead of multiple files.

Example:

COPY INTO @my_stage/output_file.csv

FROM my_table

FILE_FORMAT = (TYPE = CSV)

SINGLE = TRUE;

This is useful when a downstream process expects one output file.

Why the other options are incorrect:

A). SINGLE = TRUE does not control the file format or data type.

C). It is not an access-control setting for limiting a user's behavior.

D). It does not control the number of attempts to load or unload data.

Official Snowflake documentation reference:

Snowflake documentation for COPY INTO < location > describes SINGLE as an option that specifies whether to generate a single output file or multiple files.

Reference: Snowflake Documentation - COPY INTO < location > ; Snowflake Documentation - Unloading data from Snowflake; SnowPro Core Study Guide - Data Loading and Unloading.

**NEW QUESTION: 350**

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

**A.** ROLEADMIN

**B.** ACCOUNTADMIN

**C.** SYSADMIN

**D.** SECURITYADMIN

**E.** USERADMIN

**Answer: B,D (LEAVE A REPLY)**

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

References:

[ COF-C03 ] SnowPro Core Certification Exam Study Guide

Snowflake Documentation on Network Policies1
Section 1.3 - SnowPro Core Certification Study Guide

**NEW QUESTION: 351**
How can a user validate if a micro-partition is pruning efficiently?
**A.** By inspecting the statistics pane in the Query Profile
**B.** By inspecting the query details output in Snowsight
**C.** By using the SYSTEM$CLUSTERING_INFORMATION() function
**D.** By querying the INFORMATION_SCHEMA.TABLES view
**Answer: (SHOW ANSWER)**
The correct answer is A. By inspecting the statistics pane in the Query Profile .
The Query Profile provides execution details for a query. To evaluate micro-partition pruning, users can inspect scan statistics, including partitions scanned versus total partitions. This helps determine whether Snowflake is efficiently skipping unnecessary micro-partitions.
Why A is correct:
The Query Profile statistics pane can show how many partitions were scanned and how many total partitions existed. Efficient pruning means Snowflake scanned only a small portion of the total relevant micro-partitions.
Example interpretation:
Partitions scanned: 100
Partitions total: 10,000
This indicates strong pruning because Snowflake skipped most micro-partitions.
Why the other options are incorrect:
B). Query details in Snowsight provide useful metadata such as duration, warehouse, and query text, but the Query Profile statistics pane is the best place to validate pruning efficiency.
C). SYSTEM$CLUSTERING_INFORMATION() gives clustering information for a table, such as clustering depth and overlap, but it does not directly show whether a specific query pruned efficiently.
D). INFORMATION_SCHEMA.TABLES provides table metadata, not query-level partition pruning details.
Official Snowflake documentation reference:
Snowflake documentation explains that Query Profile provides detailed execution information, including table scan statistics such as partitions scanned and partitions total. These statistics help validate pruning effectiveness.
Reference: Snowflake Documentation - Query Profile; Snowflake Documentation - Table scan statistics; Snowflake Documentation - Micro-partitions and pruning; SnowPro Core Study Guide - Performance Optimization.
==

**NEW QUESTION: 352**
When an object is created in Snowflake. who owns the object?
**A.** The public role
**B.** The user ' s default role
**C.** The current active primary role
**D.** The owner of the parent schema
**Answer: (SHOW ANSWER)**

In Snowflake, when an object is created, it is owned by the role that is currently active. This active role is the one that is being used to execute the creation command. Ownership implies full control over the object, including the ability to grant and revoke access privileges. This is specified in Snowflake ' s documentation under the topic of Access Control, which states that " the role in use at the time of object creation becomes the owner of the object. " Snowflake Documentation: Object Ownership

**NEW QUESTION: 353**
Which Snowflake command can be used to unload the result of a query to a single file?
**A.** Use COPY INTO <external stage> followed by a GET command to download the file.
**B.** Use COPY INTO <internal stage> followed by a put command to download the file.
**C.** Use COPY INTO <internal stage> with SINGLE = TRUE followed by a GET command to download the file.
**D.** Use COPY INTO <external stage> with SINGLE = TRUE followed by a PUT command to download the file.
**Answer: (SHOW ANSWER)**
The Snowflake command to unload the result of a query to a single file is COPY INTO <internal stage> with SINGLE = TRUE followed by a GET command to download the file.This command unloads the query result into a single file in the specified internal stage

**NEW QUESTION: 354**
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: 355**
Which transformation is supported by a COPY INTO < table > command?
**A.** Filter using a where clause
**B.** Filter using a limit keyword
**C.** Cast using a SELECT statement
**D.** Order using an ORDER BY clause
**Answer: C (LEAVE A REPLY)**
The COPY INTO < table > command in Snowflake supports transformations such as casting using a SELECT statement. This allows for the transformation of data types as the data is being loaded into the table, which can be particularly useful when the data types in the source files do not match the data types in the target table

**NEW QUESTION: 356**
Which function, when combined with an INSERT statement, will insert semi-structured data into a VARIANT column?
**A.** TO_JSON
**B.** TO_VARCHAR
**C.** PARSE_JSON
**D.** OBJECT_INSERT
**Answer: (SHOW ANSWER)**
The correct answer is C. PARSE_JSON .
PARSE_JSON parses a valid JSON-formatted string and returns a VARIANT value. This is commonly used when inserting semi-structured JSON data into a VARIANT column.
Why C is correct:
A VARIANT column can store semi-structured data such as JSON objects and arrays. When the source value is a JSON string, PARSE_JSON converts it into a Snowflake VARIANT.
Example:
CREATE TABLE json_table (
src VARIANT
);
INSERT INTO json_table
SELECT PARSE_JSON( ' { " customer_id " : 101, " region " : " WEST " } ' ); Why the other options are incorrect:
A). TO_JSON converts a VARIANT value into a JSON string. It is commonly used in the opposite direction.
B). TO_VARCHAR converts a value to a string, not to a semi-structured VARIANT.
D). OBJECT_INSERT inserts or updates a key-value pair inside an existing object, but it is not the primary function used to parse JSON text into a VARIANT column.
Official Snowflake documentation reference:

Snowflake documentation describes PARSE_JSON as a function that interprets an input string as JSON and returns a VARIANT value.
Reference: Snowflake Documentation - PARSE_JSON; Snowflake Documentation - Semi-structured data types; SnowPro Core Study Guide - Working with Semi-Structured Data.
==

**NEW QUESTION: 357**
What can be used to view warehouse usage over time? (Select Two).
**A.** The load HISTORY view
**B.** The Query history view
**C.** The show warehouses command
**D.** The WAREHOUSE_METERING__HISTORY View
**E.** The billing and usage tab in the Snowflake web UI
**Answer: (SHOW ANSWER)**
To view warehouse usage over time, the Query history view and the WAREHOUSE_METERING__HISTORY View can be utilized. The Query history view allows users to monitor the performance of their queries and the load on their warehouses over a specified period1. The WAREHOUSE_METERING__HISTORY View provides detailed informationabout the workload on a warehouse within a specified date range, including average running and queued loads2. References: [COF- C03] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 358**
When semi-structured data includes dates and timestamps as string values, what Snowflake recommendation will OPTIMIZE pruning and MINIMIZE storage consumption?
**A.** Store the data as an ARRAY data type.
**B.** Flatten the key data into relational columns.
**C.** Strip the outer array using TRY_PARSE_JSON.
**D.** Store the data as a VARIANT data type.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 359**
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: 360**

Which Snowflake security feature has the HIGHEST precedence in the security feature hierarchy?

**A.** Network policies

**B.** Masking policies

**C.** Snowflake OAuth

**D.** Account roles

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 361**

Which common query problems are identified by the Query Profile? (Select TWO.)

**A.** Syntax error

**B.** Inefficient pruning

**C.** Ambiguous column names

**D.** Queries too large to fit in memory

**E.** Object does not exist or not authorized

**Answer: (SHOW ANSWER)**

The Query Profile in Snowflake can identify common query problems, including:

B). Inefficient pruning: This refers to the inability of a query to effectively limit the amount of data being scanned, potentially leading to suboptimal performance.

D). Queries too large to fit in memory: This indicates that a query requires more memory than is available in the virtual warehouse, which can lead to spilling to disk and degraded performance.

The Query Profile helps diagnose these issues by providing detailed execution statistics and visualizations, aiding in query optimization and troubleshooting.

References:

Snowflake Documentation: Query Profile

Top of Form

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

**NEW QUESTION: 362**

Which Snowflake data governance feature can support auditing when a user query reads column data?

**A.** Access History

**B.** Data classification

**C.** Column-level security

**D.** Object dependencies

**Answer: (SHOW ANSWER)**

Access History in Snowflake is a feature designed to support auditing by tracking access to data within Snowflake, including when a user ' s query reads column data. It provides detailed information on queries executed, including the user who ran the query, the query text, and the objects (e.g., tables, views) accessed by the query. This feature is instrumental for auditing purposes, helping organizations to monitor and audit data access for security and compliance.
Reference: Snowflake Documentation on Access History which details its use for auditing access to data: https://docs.snowflake.com/en/user-guide/access-history.html

**NEW QUESTION: 363**

How long is Snowpipe data load history retained?

**A.** As configured in the create pipe settings

**B.** Until the pipe is dropped

**C.** 64 days

**D.** 14 days

**Answer: (SHOW ANSWER)**

Snowpipe data load history is retained for 64 days. This retention period allows users to review and audit the data load operations performed by Snowpipe over a significant period of time, which can be crucial for troubleshooting and ensuring data integrity.
References:
[ COF-C03 ] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Snowpipe1

**NEW QUESTION: 364**

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: 365**

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: 366**
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: 367**
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: B,D (LEAVE A REPLY)**
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: 368**
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 ACCOUNTADMIH 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 propertyMINS_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: 369**
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: 370**
What does Snowflake recommend as a best practice for using secure views?
**A.** Use sequence-gen era ted values
**B.** Programmatically reveal the identifiers.
**C.** Use secure views solely for query convenience.
**D.** Do not expose the sequence-generated column(s)
**Answer: (SHOW ANSWER)**
Snowflake recommends not exposing sequence-generated columns in secure views. Secure views are used to protect sensitive data by ensuring that users can only access data for which they have permissions. Exposing sequence-generated columns can potentially reveal information about the underlying data structure or the number of rows, which might be sensitive.
Create Secure Views: Define secure views using the SECURE keyword to ensure they comply with Snowflake's security policies.
Exclude Sensitive Columns: When creating secure views, exclude columns that might expose sensitive information, such as sequence-generated columns.
CREATE SECURE VIEW secure_view AS
SELECT col1, col2
FROM sensitive_table
WHERE sensitive_column IS NOT NULL;
References:

Snowflake Documentation: Secure Views
Snowflake Documentation: Creating Secure Views
These answers and explanations should provide comprehensive guidance on the specified Snowflake topics.

**NEW QUESTION: 371**
What is the MINIMUM permission needed to access a file URL from an external stage?
**A.** MODIFY
**B.** READ
**C.** SELECT
**D.** USAGE
**Answer: (SHOW ANSWER)**
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: 372**
What Snowpark element represents a data set and provides ways to operate on the data?
**A.** List
**B.** Table
**C.** Array
**D.** DataFrame
**Answer: (SHOW ANSWER)**
The correct answer is D. DataFrame .
In Snowpark, a DataFrame represents a lazily evaluated relational data set. It provides methods to operate on data, such as selecting columns, filtering rows, joining data sets, grouping data, and applying transformations.
Why D is correct:
A Snowpark DataFrame is the primary abstraction for working with data in Snowpark. It represents data in Snowflake and lets developers build transformations using supported programming languages such as Python, Java, or Scala.
Example concept:
df = session.table( " CUSTOMERS " )
filtered_df = df.filter(df[ " COUNTRY " ] == " USA " )
The DataFrame represents the data set, and operations on it are translated into SQL that runs in Snowflake.
Why the other options are incorrect:
A). A list is a general programming data structure, not the primary Snowpark data set abstraction.
B). A table stores data in Snowflake, but in Snowpark, the object used to represent and operate on a data set programmatically is a DataFrame.
C). An array is a data structure or semi-structured type, not the main Snowpark abstraction for data operations.
Official Snowflake documentation reference:

Snowflake documentation describes Snowpark DataFrames as objects that represent a lazily evaluated relational data set and provide methods for transforming and analyzing data.
Reference: Snowflake Documentation - Snowpark DataFrames; Snowflake Documentation - Snowpark API; SnowPro Core Study Guide - SQL and Snowflake Objects.
==

**NEW QUESTION: 373**
If 3 size Small virtual warehouse is made up of two servers, how many servers make up a Large warehouse?
**A.** 4
**B.** 8
**C.** 16
**D.** 32
**Answer: (SHOW ANSWER)**
In Snowflake, each size increase in virtual warehouses doubles the number of servers. Therefore, if a sizeSmall virtual warehouse is made up of two servers, a Large warehouse, which is two sizes larger, would be made up of eight servers (2 servers for Small, 4 for Medium, and 8 for Large)2.
Size specifies the amount of compute resources available per cluster in a warehouse. Snowflake supports the following warehouse sizes:

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

snowflake
freecram.net
https://docs.snowflake.com/en/user-guide/warehouses-overview.html

**NEW QUESTION: 374**
Which command will indicate whether a materialized view is being used in a query?
**A.** DESCRIBE
**B.** EXPLAIN

**C.** SHOW
**D.** USE
**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 375**
What actions will prevent leveraging of the ResultSet cache? (Choose two.)
**A.** Removing a column from the query SELECT list
**B.** Stopping the virtual warehouse that the query is running against
**C.** Clustering of the data used by the query
**D.** Executing the RESULTS_SCAN() table function
**E.** Changing a column that is not in the cached query
**Answer: (SHOW ANSWER)**
The ResultSet cache is leveraged to quickly return results for repeated queries. Actions that prevent leveraging this cache include stopping the virtual warehouse that the query is running against (B) and executing the RESULTS_SCAN() table function (D). Stopping the warehouse clears the local diskcache, including the ResultSet cache1. The RESULTS_SCAN() function is used to retrieve the result of a previously executed query, which bypasses the need for the ResultSet cache.

**NEW QUESTION: 376**
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"
  }
}
```

freecram.net
snowflake

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.

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

**NEW QUESTION: 377**
How long can a data consumer who has a pre-signed URL access data files using Snowflake?
**A.** Indefinitely
**B.** Until the result_cache expires
**C.** Until the retention_time is met
**D.** Until the expiration time is exceeded
**Answer: C (LEAVE A REPLY)**
A data consumer who has a pre-signed URL can access data files using Snowflake until the expiration time is exceeded. The expiration time is set when the pre-signed URL is generated and determines how long the URL remains valid3.

**NEW QUESTION: 378**
Which statistics can be used to identify queries that have inefficient pruning? (Select TWO).
**A.** Bytes scanned
**B.** Bytes written to result
**C.** Partitions scanned
**D.** Partitions total
**E.** Percentage scanned from cache
**Answer: (SHOW ANSWER)**
The statistics that can be used to identify queries with inefficient pruning are 'Partitions scanned' and 'Partitions total'. These statistics indicate how much of the data was actually needed and scanned versus the total available, which can highlight inefficiencies in data pruning34.

**NEW QUESTION: 379**
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: 380**
What does Snowflake attempt to do if any of the compute resources for a virtual warehouse fail to provision during start-up?
Freecram.net
**A.** Repair the failed resources.
**B.** Restart failed resources.
**C.** Queue the failed resources
**D.** Provision the failed resources
**Answer: B (LEAVE A REPLY)**
If any compute resources for a virtual warehouse fail to provision during startup, Snowflake will attempt to restart those failed resources. The system is designed to automatically handle transient issues that might occur during the provisioning of compute resources. By restarting the failed resources, Snowflake aims to ensure that the virtual warehouse has the necessary compute capacity to handle the user ' s workloads without manual intervention.
References: Snowflake Documentation on Virtual Warehouses

**NEW QUESTION: 381**
What persistent data structures are used by the search optimization service to improve the performance of point lookups?

**A.** Micro-partitions
**B.** Clustering keys
**C.** Equality searches
**D.** Search access paths
**Answer: (SHOW ANSWER)**
The search optimization service in Snowflake uses persistent data structures known as search access paths to improve the performance of point lookups. These structures enable efficient retrieval of data by reducing the amount of data scanned during queries.
Search Access Paths:
Search access paths are special indexing structures maintained by the search optimization service.
They store metadata about the distribution of data within tables, enabling faster lookups for specific values.
Point Lookups:
Point lookups involve searching for a specific value within a column.
By leveraging search access paths, Snowflake can quickly locate the exact micro-partition containing the value, minimizing the amount of data scanned.
Performance Improvement:
The use of search access paths significantly reduces query execution time for point lookups.
This is especially beneficial for large tables where scanning all micro-partitions would be computationally expensive.
Snowflake Documentation: Search Optimization Service
Snowflake Documentation: Understanding Search Access Paths

**NEW QUESTION: 382**
A user needs to create a materialized view in the schema MYDB.MYSCHEMA.
Which statements will provide this access?
**A.** GRANT ROLE MYROLE TO USER USER1;CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO ROLE MYROLE;
**B.** GRANT ROLE MYROLE TO USER USER1;CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER USER1;
**C.** GRANT ROLE MYROLE TO USER USER1;CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER1;
**D.** GRANT ROLE MYROLE TO USER USER1;CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO MYROLE;
**Answer: (SHOW ANSWER)**
In Snowflake, to create a materialized view, the user must have the necessary privileges on the schema where the view will be created. These privileges are granted through roles, not directly to individual users.
Therefore, the correct process is to grant the role to the user and then grant the privilege to create the materialized view to the role itself.
The statement GRANT ROLE MYROLE TO USER USER1; grants the specified role to the user, allowing them to assume that role and exercise its privileges. The subsequent statement CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO MYROLE; grants the privilege to create a materialized view within the specified schema to the role MYROLE. Any user who has been granted MYROLE can then create materialized views in MYDB.MYSCHEMA.

References:
Snowflake Documentation on Roles
Snowflake Documentation on Materialized Views

**NEW QUESTION: 383**
What is run with a CALL statement?
**A.** System functions
**B.** Tasks
**C.** Stored procedures
**D.** External functions
**Answer: (SHOW ANSWER)**
The correct answer is C. Stored procedures .
In Snowflake, the CALL command is used to execute a stored procedure. Stored procedures can contain procedural logic and can be written in supported languages such as JavaScript, SQL, Python, Java, or Scala, depending on the type of procedure.
Why C is correct:
Stored procedures are invoked using the CALL statement.
Example:
CALL my_procedure();
This runs the stored procedure named my_procedure.
Why the other options are incorrect:
A). System functions are usually invoked in SQL expressions using SELECT or within statements, not with CALL.
B). Tasks are resumed, suspended, or managed using task-related commands, and they run on schedules or task graphs. They are not directly run with CALL.
D). External functions are called like regular SQL functions inside SQL statements, not with the CALL statement.
Official Snowflake documentation reference:
Snowflake documentation describes the CALL command as the command used to execute a stored procedure.
Reference: Snowflake Documentation - CALL; Snowflake Documentation - Stored procedures; SnowPro Core Study Guide - SQL and Snowflake Objects.

**NEW QUESTION: 384**
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 access3.

**NEW QUESTION: 385**
Which tool or service can be used to graphically explore and manage data stored within a Snowflake account, including databases, tables, and views?
**A.** Snowsight
**B.** Snowflake virtual warehouses
**C.** SnowCD
**D.** SnowSQL
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 386**
When using the COPY INTO < table > command, which parameter should be used to load unstructured data?
**A.** FORCE
**B.** LOAD_UNCERTAIN_FILES
**C.** FILE_PROCESSOR
**D.** LOAD_MODE
**Answer: (SHOW ANSWER)**
The FILE_PROCESSOR copy option is used when COPY INTO < table > processes unstructured files through a custom file processor. Snowflake's CUSTOM file-format type identifies a stage as containing unstructured data and is valid only with the FILE_PROCESSOR option. FORCE controls whether files already recorded as loaded are loaded again. LOAD_UNCERTAIN_FILES determines whether files whose load status is uncertain are processed, and LOAD_MODE controls how supported structured-data loading behavior is applied; neither selects an unstructured-file processor. Therefore, FILE_PROCESSOR is the correct parameter. In practice, the processor configuration determines how file content and metadata are transformed into rows and columns in the target table, while the COPY statement handles the staged-file ingestion workflow.

**NEW QUESTION: 387**
Which snowflake objects will incur both storage and cloud compute charges? (Select TWO)
**A.** Materialized view
**B.** Sequence
**C.** Secure view
**D.** Transient table
**E.** Clustered table
**Answer: (SHOW ANSWER)**
In Snowflake, both materialized views and transient tables will incur storage charges because they store data.
They will also incur compute charges when queries are run against them, as compute resources are used to process the queries. References: [COF-C03] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 388**
A table needs to be loaded. The input data is in JSON format and is a concatenation of multiple JSON documents. The file size is 3 GB. A warehouse size small is being used. The following COPY INTO command was executed:

COPY INTO SAMPLE FROM @~/SAMPLE.JSON (TYPE=JSON)
The load failed with this error:
Max LOB size (16777216) exceeded, actual size of parsed column is 17894470.
How can this issue be resolved?
**A.** Compress the file and load the compressed file.
**B.** Split the file into multiple files in the recommended size range (100 MB - 250 MB).
**C.** Use a larger-sized warehouse.
**D.** Set STRIP_OUTER_ARRAY=TRUE in the COPY INTO command.
**Answer: (SHOW ANSWER)**
The error "Max LOB size (16777216) exceeded" indicates that the size of the parsed column exceeds the maximum size allowed for a single column value in Snowflake, which is 16 MB. To resolve this issue, the file should be split into multiple smaller files that are within the recommended size range of 100 MB to 250 MB.
This will ensure that each JSON document within the files is smaller than the maximum LOB size allowed.
Compressing the file, using a larger-sized warehouse, or setting STRIP_OUTER_ARRAY=TRUE will not resolve the issue of the column size exceeding the maximum allowed. References: COPY INTO Error during Structured Data Load: "Max LOB size (16777216) exceeded..."

**NEW QUESTION: 389**
A company ' s security audit requires generating a report listing all Snowflake logins (e.g.. date and user) within the last 90 days. Which of the following statements will return the required information?
**A.** SELECT LAST_SUCCESS_LOGIN, LOGIN_NAMEFROM ACCOUNT_USAGE.USERS;
**B.** SELECT EVENT_TIMESTAMP, USER_NAMEFROM table(information_schema.
login_history_by_user())
**C.** SELECT EVENT_TIMESTAMP, USER_NAMEFROM ACCOUNT_USAGE.ACCESS_HISTORY;
**D.** SELECT EVENT_TIMESTAMP, USER_NAMEFROM ACCOUNT_USAGE.LOGIN_HISTORY;
**Answer: (SHOW ANSWER)**
To generate a report listing all Snowflake logins within the last 90 days, the ACCOUNT_USAGE.
LOGIN_HISTORY view should be used. This view provides information about login attempts, including successful and unsuccessful logins, and is suitable for security audits4.

**NEW QUESTION: 390**
Which Snowflake tasks will take advantage of underlying micro-partition metadata? Select TWO.
**A.** Query pruning
**B.** Queries using a result cache
**C.** Queries using a local disk cache
**D.** Queries using a remote disk cache
**E.** Operations using Data Manipulation Language, or DML
**Answer: (SHOW ANSWER)**
The correct answers are A. Query pruning and E. Operations using Data Manipulation Language, or DML .

Snowflake automatically gathers and maintains metadata about micro-partitions. This metadata includes information such as the range of values in columns, distinct value counts, and other statistics. Snowflake uses this metadata to optimize query execution and table operations.
Why A is correct:
Query pruning is one of the most important uses of micro-partition metadata. Snowflake can skip micro- partitions that do not contain values needed by a query predicate.
Example:
SELECT *
FROM sales
WHERE sale_date = ' 2026-01-01 ' ;
If Snowflake knows from micro-partition metadata that certain micro-partitions do not contain sale_date = ' 2026-01-01 ' , those micro-partitions can be skipped.
Why E is correct:
DML operations such as UPDATE, DELETE, and MERGE can also benefit from micro-partition metadata. Snowflake can use metadata to identify which micro-partitions are affected by the operation instead of scanning unnecessary partitions.
Why the other options are incorrect:
B). The result cache returns previously computed query results. It avoids re-executing the query and therefore does not rely on micro-partition pruning in the same way.
C). Local disk cache stores data on warehouse compute resources after access, but it is not the main feature that uses micro-partition metadata.
D). "Remote disk cache" is not the standard Snowflake concept tested here. Snowflake has remote storage, warehouse cache, and result cache, but query pruning is the metadata-driven feature.
Official Snowflake documentation reference:
Snowflake documentation explains that micro-partition metadata is used for efficient query pruning and optimization. It also supports efficient table maintenance and DML operations by helping Snowflake identify relevant micro-partitions.
Reference: Snowflake Documentation - Micro-partitions and data clustering; Snowflake Documentation - Query pruning; SnowPro Core Study Guide - Snowflake Architecture.
==

**NEW QUESTION: 391**
Which Snowflake objects track DML changes made to tables, like inserts, updates, and deletes?
**A.** Pipes
**B.** Streams
**C.** Tasks
**D.** Procedures
**Answer: B (LEAVE A REPLY)**
In Snowflake, Streams are the objects that track Data Manipulation Language (DML) changes made to tables, such as inserts, updates, and deletes. Streams record these changes along with metadata about each change, enabling actions to be taken using the changed data.This process is known as change data capture (CDC)2.

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

**NEW QUESTION: 392**
When unloading data from a query into a location, which VALIDATION_MODE parameter returns the query results?
**A.** RETURN_ < n > _ROWS
**B.** RETURN_ROWS
**C.** RETURN_ERRORS
**D.** RETURN_ALL_ERRORS
**Answer: (SHOW ANSWER)**
The correct answer is B. RETURN_ROWS .
Corrected typo:
Option D contains a typing error. RETURN_ALL_ERRORSs should be RETURN_ALL_ERRORS .
When using COPY INTO < location > to unload data, VALIDATION_MODE = RETURN_ROWS returns the query results instead of unloading them to files. This is useful for validating the result set that would be unloaded.
Why B is correct:
RETURN_ROWS is the validation mode used with unload operations to return rows generated by the query.
Example:
COPY INTO @my_stage/output/
FROM (
SELECT *
FROM customers
)
VALIDATION_MODE = RETURN_ROWS;
Why the other options are incorrect:
A). RETURN_ < n > _ROWS is associated with validation behavior for loading data, not the unload query result validation mode.
C). RETURN_ERRORS returns errors and is used for validating data load errors. It does not return query results.
D). RETURN_ALL_ERRORS returns all errors in validation mode for loading scenarios. It does not return the unload query result set.
Official Snowflake documentation reference:
Snowflake documentation for COPY INTO < location > describes VALIDATION_MODE = RETURN_ROWS as returning the rows that would be unloaded.
Reference: Snowflake Documentation - COPY INTO < location > ; Snowflake Documentation - VALIDATION_MODE; SnowPro Core Study Guide - Data Loading and Unloading.
==

**NEW QUESTION: 393**
What is a characteristic of a tag associated with a masking policy?
**A.** A tag can be dropped after a masking policy is assigned
**B.** A tag can have only one masking policy for each data type.
**C.** A tag can have multiple masking policies for each data type.
**D.** A tag can have multiple masking policies with varying data types
**Answer: (SHOW ANSWER)**
In Snowflake, a tag can be associated with only one masking policy for each data type. This means that for a given data type, you can define a single masking policy to be applied when a tag is used. Tags and masking policies are part of Snowflake ' s data classification and governance features, allowing for data masking based on the context defined by the tags.
Snowflake Documentation: Tag-Based Masking Policies

**NEW QUESTION: 394**
What happens to event notifications that are received when an automated Snowpipe is paused?
**A.** They are permanently stored in the internal metadata and processed once the pipe is resumed.
**B.** They enter a retention period of 14 days, after which they are dropped from the internal metadata.
**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: 395**
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: 396**
Which function is used to unload a relational table into a JSON file*?
**A.** PARSE_JSON

**B.** JSON_EXTRACT_PATH_TEXT
**C.** OBJECT_CONSTRUCT
**D.** TO_JSON
**Answer: (SHOW ANSWER)**
TheTO_JSONfunction in Snowflake is used to convert a relational table or individual rows into JSON format.
This function is helpful for exporting data in JSON format.
Using TO_JSON Function:
SELECT TO_JSON(OBJECT_CONSTRUCT(*))
FROM my_table;
Exporting Data:TheTO_JSONfunction converts the table rows into JSON format, which can then be exported to a file.
References:
Snowflake Documentation: TO_JSON Function
Snowflake Documentation: Exporting Data

**NEW QUESTION: 397**
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 it2.

**NEW QUESTION: 398**
Which command should be used to unload all the rows from a table into one or more files in a named stage?
**A.** COPY INTO
**B.** GET
**C.** INSERT INTO
**D.** PUT
**Answer: (SHOW ANSWER)**
To unload data from a table into one or more files in a named stage, the COPY INTO <location> command should be used. This command exports the result of a query, such as selecting all rows from a table, into files stored in the specified stage. The COPY INTO command is versatile, supporting various file formats and compression options for efficient data unloading.
References:
Snowflake Documentation: COPY INTO Location

**NEW QUESTION: 399**

Which command should be used to assign a key to a Snowflake user who needs to connect using key pair authentication?

**A.** ALTER USER jsmith SET RSA_P8_KEY= ' MIIBIjANBgkqh... ' ;

**B.** ALTER USER jsmith SET ENCRYPTED_KEY='MIIBIjANBgkqh... ' ;

**C.** ALTER USER jsmith SET RSA_PRIVATE_KEY= ' MIIBIjANBgkqh... ' ;

**D.** ALTER USER jsmith SET RSA_PUBLIC_KEY-MIIBIjANBgkqh... ' ;

**Answer: (SHOW ANSWER)**

To use key pair authentication in Snowflake, you need to set the public key for the user. This allows the user to authenticate using their private key.

Generate Key Pair:Generate a public and private key pair.

Set Public Key:

ALTER USER jsmith SET RSA_PUBLIC_KEY= ' MIIBIjANBgkqh... ' ;

Authentication:The user can now authenticate by signing requests with the corresponding private key.

References:

Snowflake Documentation: Key Pair Authentication & Key Rotation

Snowflake Documentation: ALTER USER

**NEW QUESTION: 400**

What is the MINIMUM Snowflake edition that must be used in order to see the ACCESS_HISTORY view?

Freecram.net

**A.** Standard

**B.** Enterprise

**C.** Business Critical

**D.** Virtual Private Snowflake (VPS)

**Answer: (SHOW ANSWER)**

The ACCESS_HISTORY view in Snowflake provides detailed information about queries executed in the account, including metadata such as the time of execution, the user, and the SQL text of the queries. This view is available in the Snowflake Enterprise edition and higher editions. The Standard edition does not include this feature.

Snowflake Documentation: Access History

Snowflake Editions: Snowflake Pricing

**NEW QUESTION: 401**

When unloading data: which file formal will keep floating-point number column data without truncating it?

**A.** XML

**B.** JSON

**C.** CSV

**D.** Parquet

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 402**

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

**A.** Dynamic Data Masking
**B.** Constraints
**C.** Time Travel
**D.** Row-level encryption
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 403**
Snowflake ' s access control framework combines which models for securing data? (Select TWO).
**A.** Attribute-based Access Control (ABAC 1
**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: 404**
When publishing a Snowflake Data Marketplace listing into a remote region what should be taken into consideration? (Choose two.)
**A.** There is no need to have a Snowflake account in the target region, a share will be created for each user.
**B.** The listing is replicated into all selected regions automatically, the data is not.
**C.** The user must have the ORGADMIN role available in at least one account to link accounts for replication.
**D.** Shares attached to listings in remote regions can be viewed from any account in an organization.
**E.** For a standard listing the user can wait until the first customer requests the data before replicating it to the target region.
**Answer: (SHOW ANSWER)**
When publishing a Snowflake Data Marketplace listing into a remote region, it's important to note that while the listing is replicated into all selected regions automatically, the data itself is not. Therefore, the data must be replicated separately.Additionally, the user must have the ORGADMIN role in at least one account to manage the replication of accounts1.

**NEW QUESTION: 405**
What do temporary and transient tables have m 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: 406**
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 cloningdatabase stored procedures views

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

**NEW QUESTION: 407**
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: 408**
Which Snowflake object contains all the information required to share a database?
**A.** Private listing
**B.** Secure view
**C.** Sequence
**D.** Share
**Answer: D (LEAVE A REPLY)**
In Snowflake, aShareis 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: 409**
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: 410**

Which parameter can be used to instruct a COPY command to verify data files instead of loading them into a specified table?

Freecram.net

**A.** STRIP_NULL_VALUES
**B.** SKIP_BYTE_ORDER_MARK
**C.** REPLACE_INVALID_CHARACTERS
**D.** VALIDATION_MODE
**Answer: (SHOW ANSWER)**
The VALIDATION_MODE parameter can be used with the COPY command to verify data files without loading them into the specified table.This parameter allows users to check for errors in the files

**NEW QUESTION: 411**
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 theCOPY 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: 412**
What consideration should be made when loading data into Snowflake?
**A.** Create small data files and stage them in cloud storage frequently.
**B.** Create large data files to maximize the processing overhead for each file.
**C.** The number of load operations That run in parallel can exceed the number of data files to be loaded.
**D.** The number of data files that are processed in parallel is determined by the virtual warehouse.
**Answer: (SHOW ANSWER)**
When loading data into Snowflake, one critical consideration is the parallel processing capability of the virtual warehouse used for the data loading operation. The number of data files that can be processed in parallel during a loading operation is determined by the size and resources of the virtual warehouse. A larger warehouse can process more files simultaneously, improving the efficiency and speed of data loading operations. Optimizing the size of the virtual warehouse according to the data loading needs and thesize and number of files to be loaded can significantly impact the overall performance of the data loading process.
References:
Snowflake Documentation: Optimizing Data Loading

**NEW QUESTION: 413**
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: 414**
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: 415**
When would creating and using a standard view be preferable to using a materialized view?

**A.** The underlying query of the view is run on large data sets.
**B.** The view results change often.
**C.** The underlying query of the view takes a lot of time to run.
**D.** The underlying query of the view is resource-intensive.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 416**
Which tools, services, or features should be used to write code when using the Snowpark API? (Choose two.)
**A.** Snowpark Container Services
**B.** Snowpark for Scala
**C.** Snowpark for Python
**D.** Streamlit in Snowflake
**E.** SQL Scripting
**Answer: B,C (LEAVE A REPLY)**
The Snowpark API is available through language-specific libraries for Java, Python, and Scala. Among the answer choices, Snowpark for Scala and Snowpark for Python are the two actual Snowpark libraries used to write application code with the API. Snowpark pushes supported transformations and computations into Snowflake, allowing applications to process data at scale without moving the data to the client environment. Snowpark Container Services is a managed container platform rather than a Snowpark language library. Streamlit in Snowflake is an application framework, and SQL Scripting provides procedural SQL constructs rather than a Snowpark API library. Java is also a supported Snowpark language, but it is not included in the available choices. Therefore, the two correct options are B and C.

**NEW QUESTION: 417**
How often are the Account and Table master keys automatically rotated by Snowflake?
**A.** 30 Days
**B.** 60 Days
**C.** 90 Days
**D.** 365 Days.
**Answer: (SHOW ANSWER)**
Snowflake automatically rotates the Account and Table master keys when they are more than 30 days old. Active keys are retired, and new keys are created, ensuring robust security through frequent key changes1

**NEW QUESTION: 418**
A user executed a SELECT query in Snowsight which returned a 1 GB result set. The user then downloads the files. What will occur?
**A.** The download will result in an error because the filters of the SELECT query need to be changed so that Snowsight returns a smaller result set.
**B.** The result set will be successfully downloaded from Snowsight.
**C.** The result set will be automatically compressed and the data will be downloaded as individual files.
**D.** The download will fail because the result set needs to be broken up into files no greater than 50 MB before downloading.

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 419**
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: 420**
How would you determine the size of the virtual warehouse used for a task?
**A.** Root task may be executed concurrently (i.e. multiple instances), it is recommended to leave some margins in the execution window to avoid missing instances of execution
**B.** Querying(select)the size of the stream content would help determine the warehouse size. For example, if querying large stream content, use a larger warehouse size
**C.** If using the stored procedure to execute multiple SQL statements, it ' s best to test run the stored procedure separately to size the compute resource first
**D.** Since task infrastructure is based on running the task body on schedule, it ' s recommended to configure the virtual warehouse for automatic concurrency handling using Multi-cluster warehouse (MCW) to match the task schedule
**Answer: (SHOW ANSWER)**
The size of the virtual warehouse for a task can be configured to handle concurrency automatically using a Multi-cluster warehouse (MCW). This is because tasks are designed to run their body on a schedule, and MCW allows for scaling compute resources to match the task's execution needs without manual intervention. References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 421**
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

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

**NEW QUESTION: 422**
What type of function returns one value for each Invocation?
**A.** Aggregate
**B.** Scalar
**C.** Table
**D.** Window
**Answer: (SHOW ANSWER)**
Scalar functions in Snowflake (and SQL in general) are designed to return a single value for each invocation.
They operate on a single value and return a single result, making them suitable for a wide range of data transformations and calculations within queries.
References:
Snowflake Documentation: Functions

**NEW QUESTION: 423**
Which query creates a JSON-style object with the keys name and age and the corresponding values Alice and 30?
**A.** SELECT { ' name ' : ' Alice ' , ' age ' : 30};
**B.** SELECT TO_OBJECT( ' name ' , ' Alice ' , ' age ' , 30);
**C.** SELECT OBJECT_INSERT( ' name ' , ' Alice ' , ' age ' , 30);
**D.** SELECT OBJECT_CONSTRUCT( ' name ' , ' Alice ' , ' age ' , 30);
**Answer: (SHOW ANSWER)**
OBJECT_CONSTRUCT creates a Snowflake OBJECT value from alternating key and value arguments. In option D, name is associated with Alice, and age is associated with the numeric value 30, producing a JSON-style object containing both key-value pairs. OBJECT_INSERT requires an existing OBJECT as its first argument and then adds or

updates a key-value pair, so option C does not use the required syntax. TO_OBJECT converts a compatible existing value into an OBJECT; it does not accept an alternating list of keys and values in the form shown. Snowflake also supports object constants using curly braces, but OBJECT_CONSTRUCT is the established function-based answer provided by the choices. OBJECT values can be stored directly in OBJECT columns or inside VARIANT values for semi-structured processing.

**NEW QUESTION: 424**
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)**
Thelistcommand 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: 425**
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: 426**
Which service or tool is a Command Line Interface (CLI) client used for connecting to Snowflake to execute SQL queries?
**A.** Snowsight
**B.** SnowCD
**C.** Snowpark
**D.** SnowSQL
**Answer: D (LEAVE A REPLY)**
SnowSQL is the Command Line Interface (CLI) client provided by Snowflake for executing SQL queries and performing various tasks. It allows users to connect to their Snowflake accounts and interact with the Snowflake data warehouse.
Installation: SnowSQL can be downloaded and installed on various operating systems.
Configuration: Users need to configure SnowSQL with their Snowflake account credentials.
Usage: Once configured, users can run SQL queries, manage data, and perform administrative tasks through the CLI.
References:
Snowflake Documentation: SnowSQL
Snowflake Documentation: Installing SnowSQL

**NEW QUESTION: 427**
Which statement accurately describes Snowflake ' s architecture?
**A.** It uses a local data repository for all compute nodes in the platform.
**B.** It is a blend of shared-disk and shared-everything database architectures.
**C.** It is a hybrid of traditional shared-disk and shared-nothing database architectures.
**D.** It reorganizes loaded data into internal optimized, compressed, and row-based format.
**Answer: (SHOW ANSWER)**
Snowflake ' s architecture is unique in that it combines elements of both traditional shared-disk and shared- nothing database architectures. This hybrid approach allows Snowflake to offer the scalability and performance benefits of a shared-nothing architecture (with compute and storage separated) while maintaining the simplicity and flexibility of a shared-disk architecture in managing data across all nodes in the system.
This results in an architecture that provides on-demand scalability, both vertically and horizontally, without sacrificing performance or data cohesion.

References:
Snowflake Documentation: Snowflake Architecture

**NEW QUESTION: 428**
Which Snowflake URL type allows users or applications to download or access files directly from Snowflake stage without authentication?
**A.** Directory
**B.** File
**C.** Pre-signed
**D.** Scoped
**Answer: (SHOW ANSWER)**
The pre-signed URL type allows users or applications to download or access files directly from a Snowflake stage without authentication. This URL type is open and can be used without needing to authenticate into Snowflake or pass an authorization token.

**NEW QUESTION: 429**
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-multicluster.html#:~:text=Multi%2Dcluster%20warehouses%20enable%20you,during%20peak%20and%20off%20hours.

**NEW QUESTION: 430**
When floating-point number columns are unloaded to CSV or JSON files, Snowflake truncates the values to approximately what?
Freecram.net
**A.** (12,2)
**B.** (10,4)
**C.** (14,8)
**D.** (15,9)
**Answer: (SHOW ANSWER)**
When unloading floating-point number columns to CSV or JSON files, Snowflake truncates the values to approximately 15 significant digits with 9 digits following the decimal point, which can be represented as (15,9). This ensures a balance between accuracy and efficiency in representing floating-point numbers in text- based formats, which is essential for data interchange and processing applications that consume these files.
References:

Snowflake Documentation: Data Unloading Considerations

**NEW QUESTION: 431**
What type of account can be used to share data with a consumer who does have a Snowflake account?
**A.** Data provider
**B.** Data consumer
**C.** Reader
**D.** Organization
**Answer: (SHOW ANSWER)**
A Reader account in Snowflake can be used to share data with a consumer who does not have a Snowflake account. Reader accounts are a type of shared account provided by data providers to external data consumers, allowing them to access and query shared data using Snowflake ' s web interface without needing their own Snowflake account.
References:
Snowflake Documentation: Reader Accounts

**NEW QUESTION: 432**
User INQUISITIVE_PERSON has been granted the role DATA_SCIENCE. The role DATA_SCIENCE has privileges OWNERSHIP on the schema MARKETING of the database ANALYTICS_DW.
Which command will show all privileges granted to that schema?
**A.** SHOW GRANTS ON ROLE DATA_SCIENCE
**B.** SHOW GRANTS ON SCHEMA ANALYTICS_DW.MARKETING
**C.** SHOW GRANTS TO USER INQUISITIVE_PERSON
**D.** SHOW GRANTS OF ROLE DATA_SCIENCE
**Answer: B (LEAVE A REPLY)**
To showall privileges granted to a specific schema, the command SHOW GRANTS ON SCHEMA < schema_name > should beused3. In this case, it would be SHOW GRANTS ON SCHEMA ANALYTICS_DW.MARKETING. References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 433**
How should Multi-Factor Authentication (MFA) be enabled for a user?
**A.** A USERADMIN must enroll the user
**B.** A SECURITYADMIN must enroll the user
**C.** An ACCOUNTADMIN must enroll the user
**D.** Each Snowflake user must enroll themselves
**Answer: (SHOW ANSWER)**
Users must self-enroll for MFA by logging into the web UI and scanning a QR code with an authenticator app. Administrators cannot do this on their behalf.
Reference:
Snowflake Docs: MFA Setup

**NEW QUESTION: 434**

How can a Snowflake user access near real-time metrics that can be used to identify queries affected by disk spilling?
**A.** Implement a Snowsight dashboard.
**B.** Review the history section of the Query Profile.
**C.** Use a worksheet.
**D.** Use a load monitoring chart.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 435**
Structured data must be moved from a Snowflake table to an Amazon S3 bucket and transformed into JSON.
What function should be used to do this?
**A.** PARSE_JSON()
**B.** ARRAY_CONSTRUCT()
**C.** OBJECT_CONSTRUCT()
**D.** GET
**Answer: C (LEAVE A REPLY)**

**NEW QUESTION: 436**
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: (SHOW ANSWER)**
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

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

**NEW QUESTION: 437**
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 theInformation Schema and Account Usage, the column that identifies whether a view is secure or not is IS_SECURE2.

**NEW QUESTION: 438**
Which function is used to convert rows in a relational table to a single VARIANT column?
**A.** ARRAY_AGG
**B.** OBJECT_AGG
**C.** ARRAY_CONSTRUCT
**D.** OBJECT_CONSTRUCT
**Answer: (SHOW ANSWER)**
TheOBJECT_CONSTRUCTfunction in Snowflake is used to convert rows in a relational table into a single VARIANT column that represents each row as a JSON object. This function dynamically creates a JSON object from a list of key-value pairs, where each key is a column name and each value is the corresponding column value for a row. This is particularly useful for aggregating and transforming structured data into semi- structured JSON format for further processing or analysis.
References:
Snowflake Documentation: Semi-structured Data Functions

**NEW QUESTION: 439**
How should the SPLIT_TO_TABLE( < string > , < delimiter > ) function be called?
**A.** SELECT SPLIT_TO_TABLE < COL1, ' . ' > FROM DUAL;
**B.** SELECT SPLIT_TO_TABLE( ' a.b.c ' , ' . ' );
**C.** SELECT * FROM TABLE(SPLIT_TO_TABLE( ' a.b.c ' , ' . ' ));
**D.** SELECT * FROM SPLIT_TO_TABLE( ' a-b.c ' , ' . ' );
**Answer: (SHOW ANSWER)**
SPLIT_TO_TABLE() is a table function that splits a string into multiple rows based on a delimiter. To use it correctly in SQL, you must wrap it in a TABLE() clause.
Correct Syntax:
SELECT * FROM TABLE(SPLIT_TO_TABLE( ' a.b.c ' , ' . ' ));
Reference:
Snowflake Docs: SPLIT_TO_TABLE

**NEW QUESTION: 440**
Which of the following features are available with the Snowflake Enterprise edition? (Choose two.)
**A.** Database replication and failover
**B.** Automated index management
**C.** Customer managed keys (Tri-secret secure)
**D.** Extended time travel

**E.** Native support for geospatial data

**Answer: (SHOW ANSWER)**

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

**NEW QUESTION: 441**

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: 442**

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: 443**
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 become temporarily inactive until they are manually reactivated.
**D.** The tags are automatically cloned with the schema, maintaining all associations.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 444**
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: 445**

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: D (LEAVE A REPLY)**
The primary factor that determines the cost of using a virtual warehouse in Snowflake isD. 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: 446**
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: 447**
What is the purpose of the Snowflake SPLIT TO_TABLE function?
**A.** To count the number of characters in a string
**B.** To split a string into an array of sub-strings
**C.** To split a string and flatten the results into rows
**D.** To split a string and flatten the results into columns
**Answer: (SHOW ANSWER)**

The purpose of the Snowflake SPLIT_TO_TABLE function is to split a string based on a specified delimiter and flatten the results into rows. This table function is useful for transforming a delimited string into a set of rows that can be further processed or queried5.

**NEW QUESTION: 448**
Which role must be used to create resource monitors?
**A.** SECURITYADMIN
**B.** ACCOUNTADMIN
**C.** SYSADMIN
**D.** ORGADMIN
**Answer: (SHOW ANSWER)**
In Snowflake, the ACCOUNTADMIN role is required to create resource monitors. Resource monitors are used to manage and monitor the consumption of compute resources. The ACCOUNTADMIN role has the necessary privileges to create, configure, and manage resource monitors across the account.
Snowflake Documentation: Resource Monitors

**NEW QUESTION: 449**
Which metadata table will store the storage utilization information even for dropped tables?
**A.** DATABASE_STORAGE_USAGE_HISTORY
**B.** TABLE_STORAGE_METRICS
**C.** STORAGE_DAILY_HISTORY
**D.** STAGE STORAGE USAGE HISTORY
**Answer: B (LEAVE A REPLY)**
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: 450**
Which privilege grants the ability to set a column-level security masking policy on a table or view column?
Freecram.net
**A.** APPLY
**B.** CREATE
**C.** SET
**D.** MODIFY
**Answer: C (LEAVE A REPLY)**
In Snowflake, theAPPLYprivilege 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 theCREATE MASKING POLICYcommand.
Grant APPLY Privilege:Grant theAPPLYprivilege on the masking policy to the relevant roles or users.
Apply the Masking Policy:Use theALTER TABLEorALTER VIEWcommand 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: 451**
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.

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

**NEW QUESTION: 452**
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: 453**
What Snowflake features are recommended to restrict unauthorized users from accessing Personal Identifiable Information (PII)? (Select TWO).
**A.** Secure views
**B.** Dynamic Data Masking
**C.** Transient tables
**D.** Multi-Factor Authentication (MFA)

**E.** Data encryption
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 454**
What should be considered before implementing a Snowflake Data Clean Room?
**A.** Participants can run unrestricted SQL queries directly against another participant's protected data.
**B.** Providers can configure and share clean rooms only with individual users in other Snowflake accounts.
**C.** A consumer must be approved by a clean-room administrator before the consumer can be invited.
**D.** The regional locations of the participating Snowflake accounts must be considered because collaboration availability and cross-region requirements depend on the clean-room deployment.
**Answer: (SHOW ANSWER)**
Snowflake Data Clean Rooms provide controlled collaboration between parties without exposing the underlying protected data directly. Participants perform approved analyses through the clean room's permitted templates and policies; they do not receive unrestricted access to another party's raw data.
The Snowflake account locations of the provider and consumer are an implementation consideration. Earlier Data Clean Room implementations commonly required collaborators to use accounts in the same supported region. Newer Snowflake capabilities may support cross-region or cross-cloud collaboration under additional configuration and availability requirements.
Therefore, this question should not state an unconditional, permanent rule that every clean-room participant must always be in the same geographic region. The technically corrected option is that region compatibility and supported cross-region capabilities must be evaluated before implementation.
The original version of this question likely expected "All participating accounts must be in the same geographic region." That answer should be used only when the applicable SnowPro study guide or exam version explicitly documents that limitation.
References:
Snowflake Documentation - Introduction to Snowflake Data Clean Rooms: https://docs.snowflake.com/en/user-guide/cleanrooms/introduction
Snowflake Documentation - Data Clean Room provider and consumer workflows: https://docs.snowflake.com/en/user-guide/cleanrooms/working-with-cleanrooms
==

**NEW QUESTION: 455**
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: 456**
Which Query Profile metrics can be used to determine if a query is pruning efficiently?
**A.** Bytes spilled to local storage and Bytes spilled to remote storage
**B.** Number of rows inserted and Number of rows updated
**C.** Partitions scanned and Partitions total
**D.** Bytes read from result and Bytes written to result
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 457**
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: (SHOW ANSWER)**
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.

**NEW QUESTION: 458**
Data storage for individual tables can be monitored using which commands and/or objects? (Choose two.)
**A.** SHOW STORAGE BY TABLE;
**B.** SHOW TABLES;
**C.** Information Schema -> TABLE_HISTORY
**D.** Information Schema -> TABLE_FUNCTION
**E.** Information Schema -> TABLE_STORAGE_METRICS
**Answer: (SHOW ANSWER)**
To monitor data storage for individual tables, the commands and objects that can be used are 'SHOW STORAGE BY TABLE;' and the Information Schema view 'TABLE_STORAGE_METRICS'.These tools provide detailed information about the storage utilization for tables. References: Snowflake Documentation

**NEW QUESTION: 459**
How are serverless features billed?
**A.** Per second multiplied by an automatic sizing for the job
**B.** Per minute multiplied by an automatic sizing for the job, with a minimum of one minute
**C.** Per second multiplied by the size, as determined by the SERVERLESS_FEATURES_SIZE account parameter

**D.** Serverless features are not billed, unless the total cost for the month exceeds 10% of the warehouse credits, on the account

**Answer: (SHOW ANSWER)**

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

**NEW QUESTION: 460**

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.

==

**NEW QUESTION: 461**

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 tables3.

**NEW QUESTION: 462**

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

**A.** Configure the maximum clusters setting to " Maximum. "

**B.** Turn on the additional clusters manually after starting the warehouse.

**C.** Set the minimum Clusters and maximum Clusters settings to the same value.

**D.** Set the minimum clusters and maximum clusters settings to different values.

**Answer: (SHOW ANSWER)**

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

**NEW QUESTION: 463**

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: 464**
What is a key difference between JSON and CSV file formats when loading data? Select TWO.
**A.** CSV stores flat, tabular data
**B.** CSV supports complex data types such as arrays and objects
**C.** JSON must be converted to CSV before it can be queried
**D.** CSV files automatically detect and apply data types during loading
**E.** JSON can be loaded into a VARIANT column without predefined schema
**Answer: (SHOW ANSWER)**
The correct answers are A. CSV stores flat, tabular data and E. JSON can be loaded into a VARIANT column without predefined schema .
Important correction:
The provided answer A, B is not correct. CSV does not natively support complex semi-structured data types such as arrays and objects. JSON does.
Why A is correct:
CSV is a delimited, flat file format commonly used for structured, tabular data. Each row usually represents a record, and each field is separated by a delimiter such as a comma.
Why E is correct:
JSON is semi-structured data and can be loaded into a Snowflake VARIANT column. This allows Snowflake to store JSON objects and arrays without requiring every nested field to be predefined as a relational column.
Example:
```
CREATE TABLE json_data (
src VARIANT
);
COPY INTO json_data
FROM @my_stage
FILE_FORMAT = (TYPE = JSON);
```
Why the other options are incorrect:
B). CSV does not natively support arrays and objects. JSON supports nested structures such as arrays and objects.

C). JSON does not need to be converted to CSV before it can be queried in Snowflake. Snowflake can query JSON data stored in VARIANT, OBJECT, or ARRAY columns.
D). CSV files do not automatically detect and apply Snowflake table data types during loading. The target table schema and file format options control how data is interpreted.
Official Snowflake documentation reference:
Snowflake documentation explains that CSV is used for structured data, while semi-structured formats such as JSON can be loaded into VARIANT columns and queried using semi-structured data syntax.
Reference: Snowflake Documentation - Semi-structured data; Snowflake Documentation - Loading JSON data; Snowflake Documentation - CSV file format options; SnowPro Core Study Guide - Data Loading and Unloading.
==

**NEW QUESTION: 465**
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: 466**
Use of which feature or setting consumes Snowflake credits, but generates no storage costs?
**A.** Materialized view
**B.** Search optimization service
**C.** Query acceleration service
**D.** Table clustering
**Answer: (SHOW ANSWER)**

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

**NEW QUESTION: 467**
What is a best practice after creating a custom role?

**A.** Create the custom role using the SYSADMIN role.
**B.** Assign the custom role to the SYSADMIN role
**C.** Assign the custom role to the PUBLIC role
**D.** Add__CUSTOM to all custom role names
**Answer: (SHOW ANSWER)**
Assigning the custom role to the SYSADMIN role is considered a best practice because it allows the SYSADMIN role to manage objects created by the custom role. This is important for maintaining proper access control and ensuring that the SYSADMIN can perform necessary administrative tasks on objects created by users with the custom role.
References:
[ COF-C03 ] SnowPro Core Certification Exam Study Guide
Section 1.3 - SnowPro Core Certification Study Guide1

**NEW QUESTION: 468**
Given the statement template below, which database objects can be added to a share?(Select TWO).
GRANT < privilege > ON < object > < object_name > To SHARE < share_name > ;
**A.** Secure functions
**B.** Stored procedures
**C.** Streams
**D.** Tables
**E.** Tasks
**Answer: C,D (LEAVE A REPLY)**
In Snowflake, shares are used to share data across different Snowflake accounts securely. When you create a share, you can include various database objects that you want to share with consumers. According to Snowflake ' s documentation, the types of objects that can be shared include tables, secure views, secure materialized views, and streams. Secure functions and stored procedures are not shareable objects. Tasks also cannot be shared directly. Therefore, the correct answers are streams (C) and tables (D).
To share a stream or a table, you use theGRANTstatement to grant privileges on these objects to a share. The syntax for sharing a table or stream involves specifying the type of object, the object name, and the share to which you are granting access. For example:
GRANTSELECTONTABLEmy_tableTOSHARE my_share;GRANTSELECTONSTREAM my_streamTOSHARE my_share;
These commands grant the SELECT privilege on a table namedmy_tableand a stream namedmy_streamto a share namedmy_share. This enables the consumer of the share to access these objects according to the granted privileges.
Reference: Snowflake Documentation on Shares and Database Objects (https://docs.snowflake.com)

**NEW QUESTION: 469**
Which default warehouse configuration has the highest precedence whenever a new session is created by a user?
**A.** Default warehouse tor the user
**B.** Default warehouse of the role assigned to the user
**C.** Default warehouse in the configuration file of the client utilities
**D.** Default warehouse specified on a CLI or in drivers/connectors parameters

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 470**

By default, when is a resource monitor reset to 0?

**A.** At the beginning of each day

**B.** At the beginning of each week

**C.** At the beginning of each month

**D.** At the beginning of each year

**Answer: (SHOW ANSWER)**

The correct answer is C. At the beginning of each month .

A Snowflake resource monitor tracks credit usage for an account or specific warehouses. Resource monitors can have reset frequencies such as daily, weekly, monthly, yearly, or never. By default, the reset frequency is monthly.

Why C is correct:

With the default monthly frequency, the resource monitor's credit usage resets to 0 at the beginning of each monthly interval.

Why the other options are incorrect:

A). Daily reset can be configured, but it is not the default.

B). Weekly reset can be configured, but it is not the default.

D). Yearly reset can be configured, but it is not the default.

Official Snowflake documentation reference:

Snowflake documentation explains that resource monitors support configurable reset frequencies and that the default frequency is monthly.

Reference: Snowflake Documentation - Resource monitors; Snowflake Documentation - Credit quota and monitor frequency; SnowPro Core Study Guide - Snowflake Account and Resource Management.

**NEW QUESTION: 471**

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

**A.** Use public and private key pair authentication

**B.** Specify using SNOWSQL_PWD environment variables

**C.** Run a web-based authorization flow

**D.** Enter through an interactive prompt

**E.** Use an OAuth token

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 472**

Which Snowflake tool would be BEST to troubleshoot network connectivity?

**A.** SnowCLI

**B.** SnowUI

**C.** SnowSQL

**D.** SnowCD

**Answer: D (LEAVE A REPLY)**

SnowCD (Snowflake Connectivity Diagnostic Tool) is the best tool provided by Snowflake for troubleshooting network connectivity issues. It helps diagnose and resolve issues related to connecting to Snowflake services https://docs.snowflake.com/en/user-guide/snowcd.html#:~:text=SnowCD%20(i.e.%20Snowflake%20Connectivity%20Diagnostic,their%20network%20connection%20to%20Snowflake.

**NEW QUESTION: 473**
What is the minimum Snowflake edition that has column-level security enabled?
**A.** Standard
**B.** Enterprise
**C.** Business Critical
**D.** Virtual Private Snowflake
**Answer: (SHOW ANSWER)**
Column-level security, which allows for the application of masking policies to columns in tables or views, is available starting from the Enterprise edition of Snowflake1.
References = [COF-C03] SnowPro Core Certification Exam Study Guide, Snowflake Documentation1

**NEW QUESTION: 474**
A Snowflake user wants to optimize performance for a query that queries only a small number of rows in a table. The rows require significant processing. The data in the table does not change frequently.
What should the user do?
**A.** Add a clustering key to the table.
**B.** Add the search optimization service to the table.
**C.** Create a materialized view based on the query.
**D.** Enable the query acceleration service for the virtual warehouse.
**Answer: C (LEAVE A REPLY)**
In a scenario where a Snowflake user queries only a small number of rows that require significant processing and the data in the table does not change frequently, the most effective way to optimize performance is by creating a materialized view based on the query. Materialized views store the result of the query and can significantly reduce the computation time for queries that are executed frequently over unchanged data.
Why Materialized Views:Materialized views precompute and store the result of the query. This is especially beneficial for queries that require heavy processing. Since the data does not change frequently, the materialized view will not need to be refreshed often, making it an ideal solution for this use case.
Implementation Steps:
To create a materialized view, use the following SQL command:
CREATEMATERIALIZEDVIEWmy_materialized_viewASSELECT...FROMmy_tableWHERE...; When the query is run, Snowflake uses the precomputed results from the materialized view, thus skipping the need for recalculating the data and improving query performance.
Reference:For more information on materialized views and how they can be used to optimize query performance, refer to the Snowflake documentation on materialized views: https://docs.snowflake.com/en/user-guide/views-materialized.html

**NEW QUESTION: 475**
What object does Snowflake recommend using when planning to unload similarly-formatted data on a regular basis?
**A.** Stream
**B.** Storage integration
**C.** Named file format
**D.** Task
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 476**
Which clients does Snowflake support Multi-Factor Authentication (MFA) token caching for? (Select TWO).
**A.** GO driver
**B.** Node.js driver
**C.** ODBC driver
**D.** Python connector
**E.** Spark connector
**Answer: (SHOW ANSWER)**
Multi-Factor Authentication (MFA) token caching is typically supported for clients that maintain a persistent connection or session with Snowflake, such as the ODBC driver and Python connector, to reduce the need for repeated MFA challenges.References:Based on general security practices in cloud services as of 2021.

**NEW QUESTION: 477**
A global company needs to share financial reports with managers who are in different regions. Managers should only see the data from their respective regions. How can these requirements be met with the least amount of operational overhead?
**A.** Create unique views for each region.
**B.** Build individual reports for each manager.
**C.** Use row-level security on the source data table used to generate the reports.
**D.** Apply Dynamic Data Masking based on a custom role assigned to each manager.
**Answer: (SHOW ANSWER)**
The correct answer is C. Use row-level security on the source data table used to generate the reports .
The requirement is to let managers see only rows for their respective regions. This is a row-filtering requirement, so Snowflake row-level security is the best fit.
Why C is correct:
Row-level security in Snowflake is implemented using row access policies. A row access policy can dynamically filter table rows based on the current user, current role, or mapping table logic.
Example concept:
CREATE ROW ACCESS POLICY region_policy
AS (region STRING) RETURNS BOOLEAN - >
region IN (
SELECT allowed_region
FROM manager_region_map

WHERE manager_role = CURRENT_ROLE()
);
This allows one shared table or report to serve multiple managers while enforcing region-specific visibility.
Why the other options are incorrect:
A). Creating unique views for each region increases operational overhead as regions or access rules change.
B). Building individual reports for each manager creates the highest operational overhead and is difficult to maintain.
D). Dynamic Data Masking hides or transforms column values. It does not filter rows by region.
Official Snowflake documentation reference:
Snowflake documentation describes row access policies as schema-level objects that determine whether a row is visible in a query result. They are used to implement row-level security.
Reference: Snowflake Documentation - Row access policies; Snowflake Documentation - Dynamic Data Masking; SnowPro Core Study Guide - Data Protection and Governance.

**NEW QUESTION: 478**
What happens when a network policy includes values that appear in both the allowed and blocked IP address list?
**A.** Those IP addresses are allowed access to the Snowflake account as Snowflake applies the allowed IP address list first.
**B.** Those IP addresses are denied access lei the Snowflake account as Snowflake applies the blocked IP address list first.
**C.** Snowflake issues an alert message and adds the duplicate IP address values lo both 'he allowed and blocked IP address lists.
**D.** Snowflake issues an error message and adds the duplicate IP address values to both the allowed and blocked IP address list
**Answer: (SHOW ANSWER)**
In Snowflake, when setting up a network policy that specifies both allowed and blocked IP address lists, if an IP address appears in both lists, access from that IP address will be denied. The reason is that Snowflake prioritizes security, and the presence of an IP address in the blocked list indicates it should not be allowed regardless of its presence in the allowed list. This ensures that access controls remain stringent and that any potentially unsafe IP addresses are not inadvertently permitted access.
References:
Snowflake Documentation: Network Policies

**NEW QUESTION: 479**
What is a characteristic of materialized views in Snowflake?
**A.** Materialized views do not allow joins.
**B.** Clones of materialized views can be created directly by the user.
**C.** Multiple tables can be joined in the underlying query of a materialized view.
**D.** Aggregate functions can be used as window functions in materialized views.
**Answer: C (LEAVE A REPLY)**
One of the characteristics of materialized views in Snowflake is that they allow multiple tables to be joined in the underlying query. This enables the pre-computation of complex queries involving joins, which can significantly improve

the performance of subsequent queries that access the materialized view4. References: [COF-C03] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 480**

What is the EXPLAIN function used for in Snowflake?

**A.** To compile a SQL statement and provide an execution plan using a virtual warehouse

**B.** To compile a SQL statement and provide an execution plan while consuming cloud service credits

**C.** To execute a SQL statement and provide a history of the query execution and the number of compute credits consumed

**D.** To execute a SQL statement and provide a history of the query execution, including the number of cloud services credits consumed

**Answer: B (LEAVE A REPLY)**

The correct answer is B. To compile a SQL statement and provide an execution plan while consuming cloud service credits .

The EXPLAIN command returns the logical execution plan for a SQL statement. It compiles the SQL statement but does not execute it. Because the statement is compiled and optimized by Snowflake's cloud services layer, it does not require a running virtual warehouse.

Why B is correct:

EXPLAIN generates an execution plan without running the query. The compile and optimization work happens in the cloud services layer. Cloud services usage may apply according to Snowflake's normal cloud services credit rules.

Why the other options are incorrect:

A). EXPLAIN does not require a virtual warehouse to execute the query because the query is not actually run.

C). EXPLAIN does not execute the SQL statement and does not provide actual query history or compute credit consumption.

D). EXPLAIN does not execute the statement or return execution history.

Official Snowflake documentation reference:

Snowflake documentation describes EXPLAIN as a command that returns the logical execution plan for a SQL statement without executing the statement. Snowflake's cloud services layer handles query parsing and optimization.

Reference: Snowflake Documentation - EXPLAIN; Snowflake Documentation - Snowflake architecture cloud services layer; SnowPro Core Study Guide - Performance Optimization.

==

**NEW QUESTION: 481**

Which command can be used to load data into an internal stage?

**A.** LOAD

**B.** copy

**C.** GET

**D.** PUT

**Answer: (SHOW ANSWER)**

The PUT command is used to load data into an internal stage in Snowflake. This command uploads data files from a local file system to a named internal stage, making the data available for subsequent loading into a Snowflake table using the COPY INTO command.
References:
[ COF-C03 ] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Data Loading

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

**NEW QUESTION: 482**
What causes objects in a data share to become unavailable to a consumer account?
**A.** The DATA_RETENTI0N_TIME_IN_DAYS parameter in the consumer account is set to 0.
**B.** The consumer account runs the GRANT IMPORTED PRIVILEGES command on the data share every 24 hours.
**C.** The objects in the data share are being deleted and the grant pattern is not re-applied systematically.
**D.** The consumer account acquires the data share through a private data exchange.
**Answer: C (LEAVE A REPLY)**
Objects in a data share become unavailable to a consumer account if the objects in the data share are deleted or if the permissions on these objects are altered without re-applying the grant permissions systematically.
This is because the sharing mechanism in Snowflake relies on explicit grants of permissions on specific objects (like tables, views, or secure views) to the share. If these objects are deleted or if their permissions change without updating the share accordingly, consumers can lose access.
TheDATA_RETENTION_TIME_IN_DAYSparameter does not directly affect the availability of shared objects, as it controls how long Snowflake retains historical data for time travel and does not impact data sharing permissions.
Running theGRANT IMPORTED PRIVILEGEScommand in the consumer account is not related to the availability of shared objects; this command is used to grant privileges on imported objects within the consumer ' s account and is not a routine maintenance command that would need to be run regularly.
Acquiring a data share through a private data exchange does not inherently make objects unavailable; issues would only arise if there were problems with the share configuration or if the shared objects were deleted or had their permissions altered without re-granting access to the share.
Reference: Snowflake Documentation on Managing Access to Shared Data (https://docs.snowflake.com/en /user-guide/data-sharing-consuming.html#managing-access-to-shared-data)

**NEW QUESTION: 483**
A column named " Data " contains VARIANT data and stores values as follows:

How will Snowflake extract the employee ' s name from the column data?

**A.** Data:employee.name

**B.** DATA:employee.name

**C.** data:Employee.name

**D.** data:employee.name

**Answer: (SHOW ANSWER)**

In Snowflake, to extract a specific value from a VARIANT column, you use the column name followed by a colon and then the key. The keys are case-sensitive. Therefore, to extract the employee's name from the "Data" column, the correct syntax is data:employee.name.

**NEW QUESTION: 484**

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: 485**

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

**A.** SCIM

**B.** Federated authentication

**C.** TLS 1.2

**D.** Key-pair authentication

**E.** OAuth

**F.** OCSP authentication

**Answer: (SHOW ANSWER)**

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

**NEW QUESTION: 486**

Which table type is used in the file processing pipeline to process unstructured data in Snowflake?

**A.** Transient

**B.** Standard
**C.** Temporary
**D.** Directory
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 487**
What are the main differences between the account usage views and the information schema views? (Select TWO).
**A.** No active warehouse to needed to query account usage views but one is needed to query information schema views.
**B.** Account usage views do not contain data about tables but information schema views do.
**C.** Account issue views contain dropped objects but information schema views do not.
**D.** Data retention for account usage views is 1 year but is 7 days to 6 months for information schema views, depending on the view.
**E.** Information schema views are read-only but account usage views are not.
**Answer: (SHOW ANSWER)**
The account usage views in Snowflake provide historical usage data about the Snowflake account, and they retain this data for a period of up to 1 year. These views include information about dropped objects, enabling audit and tracking activities. On the other hand, information schema views provide metadata about database objects currently in use, such as tables and views, but do not include dropped objects. The retention of data in information schema views varies, but it is generally shorter than the retention for account usage views, ranging from 7 days to a maximum of 6 months, depending on the specific view.
References: Snowflake Documentation on Account Usage and Information Schema

**NEW QUESTION: 488**
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: 489**
What is the minimum Snowflake edition needed for database failover and fail-back between Snowflake accounts for business continuity and disaster recovery?
**A.** Standard
**B.** Enterprise
**C.** Business Critical
**D.** Virtual Private Snowflake

**Answer: (SHOW ANSWER)**

The minimum Snowflake edition required for database failover and fail-back between Snowflake accounts for business continuity and disaster recovery is the Business Critical edition. References: Snowflake Documentation3.

**NEW QUESTION: 490**

What does the Remote Disk I/O statistic in the Query Profile indicate?

**A.** Time spent reading from the result cache.

**B.** Time spent reading from the virtual warehouse cache.

**C.** Time when the query processing was blocked by remote disk access.

**D.** The level of network activity between the Cloud Services layer and the virtual warehouse.

**Answer: (SHOW ANSWER)**

The Remote Disk I/O statistic in the Query Profile reflects time spent waiting on remote disk access, which can occur when data needs to be retrieved from external storage (remote). This metric is crucial for identifying bottlenecks related to I/O delays, often suggesting a need for performance optimization in data retrieval paths.

The other options relate to caching and network activity, but Remote Disk I/O specifically measures the wait time for data access from remote storage locations.

**NEW QUESTION: 491**

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

**A.** UTF-8

**B.** UTF-16

**C.** ISO S859-1

**D.** ANSI_X3.A

**Answer: A (LEAVE A REPLY)**

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: 492**

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: 493**
How can customer data be restricted by country for a global company using a multi-tenant environment?
**A.** Implement a new account structure and create sub-accounts for each country. Then replicate the required data.
**B.** Implement a secure User-Defined Function (UDF) and use a secure data share for each country.
**C.** 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.
**D.** Implement a column-level security policy in the customer table that restricts access based on the user ' s geographic location.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 494**
Which schema-level objects allow the user to periodically perform an action under specific conditions, based on data within Snowflake?
**A.** Materialized views
**B.** Secure views
**C.** Alerts
**D.** External tables
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 495**
Which strategy is recommended to improve the performance of a large, complex query?
**A.** Add more storage
**B.** Convert to a multi-cluster warehouse
**C.** Increase the size of the warehouse
**D.** Run the query during off-peak hours
**Answer: (SHOW ANSWER)**
To improve the performance of compute-intensive queries, increasing the warehouse size (e.g., from Medium to Large or X-Large) provides more compute resources (CPU/memory) , allowing Snowflake to process queries faster. This is the most direct and effective method for performance optimization.
Reference:
Snowflake Docs: Warehouse Sizing
"Larger warehouses provide more compute resources for faster performance."

**NEW QUESTION: 496**
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.

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

**NEW QUESTION: 497**
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: 498**
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: 499**
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: (SHOW ANSWER)**
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: 500**
In which use cases does Snowflake apply egress charges?
**A.** Data sharing within a specific region
**B.** Query result retrieval
**C.** Database replication
**D.** Loading data into Snowflake
**Answer: (SHOW ANSWER)**
Snowflake applies egress charges in the case of database replication when data is transferred out of a Snowflake region to another region or cloud provider. This is because the data transfer incurs costs associated with moving data across different networks. Egress charges are not applied for data sharing within the same region, query result retrieval, or loading data into Snowflake, as these actions do not involve data transfer across regions.
References:
[COF-C03] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Data Replication and Egress Charges1

**NEW QUESTION: 501**
What is a Snowflake recommendation for the use of key pair authentication?
**A.** Keys should be stored in a secure location on the ACCOUNTADMIN file system.
**B.** Operator access should be limited to read-only access.
**C.** Key pair rotation should be automated.
**D.** Role access should be designed to cover the operations that the user account anticipates needing.
**Answer: (SHOW ANSWER)**
Snowflake recommends automating key pair rotation to improve security and comply with industry best practices. SSH key pairs are used for key-pair authentication with clients like SnowSQL or Snowpipe.
Reference:
Snowflake Docs: Key Pair Authentication & Management

**NEW QUESTION: 502**
A Snowflake practitioner wants to improve the performance of a pipeline that repeatedly queries semi- structured data.
Which features can improve query performance for this workload? Select TWO.
**A.** Materialized view
**B.** A clustering key directly on the semi-structured column
**C.** Query Acceleration Service
**D.** Search Optimization Service

**E.** Multi-cluster virtual warehouse

**Answer: (SHOW ANSWER)**

The Search Optimization Service can improve highly selective searches against supported values stored in semi-structured columns such as VARIANT, OBJECT, and ARRAY. It maintains a search-access structure that helps Snowflake locate qualifying micro-partitions without scanning all candidate data.

A materialized view can precompute frequently used expressions or extracted values from semi-structured data. Repeated queries can then access the materialized representation instead of repeatedly performing expensive extraction or transformation operations against the full source table.

A clustering key cannot simply be defined directly on a VARIANT, OBJECT, or ARRAY value. In some designs, a supported expression that extracts and converts a scalar value can be used for clustering, but option B as written is not generally correct.

Query Acceleration Service can help certain eligible scan-intensive queries, but it is not the semi-structured- data-specific optimization represented by this question. Multi-cluster warehouses improve concurrency rather than the execution speed of one query.

References:

Snowflake Documentation - Search optimization for semi-structured queries: https://docs.snowflake.com/en /user-guide/search-optimization/semi-structured-queries

Snowflake Documentation - Materialized views: https://docs.snowflake.com/en/user-guide/views- materialized

==

**NEW QUESTION: 503**

What is the MINIMUM Snowflake edition required to use a dedicated metadata store?

**A.** Standard

**B.** Business Critical

**C.** Virtual Private Snowflake (VPS)

**D.** Enterprise

**Answer: B (LEAVE A REPLY)**

**NEW QUESTION: 504**

What are type predicates used for?

**A.** Extracting data from a variant column

**B.** Casting a value in a variant column to a particular data type

**C.** Determining if a value in a variant column is a particular data type

**D.** Manipulating objects and arrays in a VARIANT column

**Answer: C (LEAVE A REPLY)**

Type predicates in Snowflake are used to determine if a value in a VARIANT column is of a particular data type. This is useful when working with semi-structured data stored in VARIANT columns, as it allows for data type validation and conditional processing based on the data type.

Snowflake Documentation: Type Predicates

**NEW QUESTION: 505**

What action should be taken if a large number of concurrent queries are queued in a virtual warehouse?
**A.** Disable auto-suspend on the warehouse.
**B.** Scale-up by resizing the warehouse.
**C.** Scale-out with a multi-cluster warehouse.
**D.** Enable auto-resume on the warehouse.
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 506**
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: 507**
What is a non-configurable feature that provides historical data that Snowflake may recover during a 7-day period?
**A.** Fail-safe
**B.** Time Travel
**C.** Cloning
**D.** Account replication
**Answer: (SHOW ANSWER)**
Fail-safe is a non-configurable feature in Snowflake that provides an additional layer of data protection beyond Time Travel. Time Travel allows users to access historical data within a configurable period (up to 90 days), while Fail-safe provides an additional 7-day period during which Snowflake retains historical data to recover from significant data loss or corruption incidents. This period is not accessible by users but can be used by Snowflake support to assist in data recovery efforts.
References: Snowflake Documentation on Fail-safe and Time Travel

**NEW QUESTION: 508**
Which command is used to unload data from a Snowflake database table into one or more files in a Snowflake stage?
**A.** CREATE STAGE
**B.** COPY INTO < table >
**C.** COPY INTO < location >

**D.** CREATE PIPE

**Answer: (SHOW ANSWER)**

The COPY INTO < location > command is used to unload data from a Snowflake database table into one or more files in a Snowflake stage1.

**NEW QUESTION: 509**

A size X-Small virtual warehouse ran for 90 seconds and was shut down. The warehouse was then run for another 30 seconds before being shut down again.

How many seconds will be billed?

**A.** 120 seconds

**B.** 90 seconds

**C.** 150 seconds

**D.** 180 seconds

**Answer: (SHOW ANSWER)**

**NEW QUESTION: 510**

In the Data Exchange, who can get or request data from the listings? (Select TWO).

**A.** Users with ACCOUNTADMIN role

**B.** Users with sysadmin role

**C.** Users with ORGADMIN role

**D.** Users with import share privilege

**E.** Users with manage grants privilege

**Answer: (SHOW ANSWER)**

In the Snowflake Data Exchange, the ability to get or request data from listings is generally controlled by specific roles and privileges:

A . Users with ACCOUNTADMIN role: This role typically has the highest level of access within a Snowflake account, including the ability to manage and access all features and functions. Users with this role can access data listings within the Data Exchange.

D . Users with import share privilege: This specific privilege is necessary for users who need to import shared data from the Data Exchange. This privilege allows them to request and access data listings explicitly shared with them.

**NEW QUESTION: 511**

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)**

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

**NEW QUESTION: 512**
What is the MAXIMUM Time Travel retention period for a transient table?
**A.** O days
**B.** 1 day
**C.** 7 days
**D.** 90 days
**Answer: (SHOW ANSWER)**
The maximum Time Travel retention period for a transient table in Snowflake is 1 day.This is the default and maximum duration for which Snowflake maintains the historical data for transient tables, allowing users to query data as it appeared at any point within the past 24 hours2.

**NEW QUESTION: 513**
What impacts the credit consumption of maintaining a materialized view? (Choose two.)
**A.** Whether or not it is also a secure view
**B.** How often the underlying base table is queried
**C.** How often the base table changes
**D.** Whether the materialized view has a cluster key defined
**E.** How often the materialized view is queried
**Answer: (SHOW ANSWER)**
The credit consumption for maintaining a materialized view is impacted by how often the base table changes and whether the materialized view has a cluster key defined (D). Changes to the base table can trigger a refresh of the materialized view, consuming credits. Additionally, having a cluster key defined can optimize the performance and credit usage during the materialized view's maintenance. References: SnowPro Core Certification materialized viewcredit consumption

**NEW QUESTION: 514**
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: (SHOW ANSWER)**
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: 515**
How does Snowflake reorganize data when it is loaded? (Select TWO).
**A.** Binary format
**B.** Columnar format
**C.** Compressed format
**D.** Raw format
**E.** Zipped format
**Answer: (SHOW ANSWER)**
When data is loaded into Snowflake, it undergoes a reorganization process where the data is stored in a columnar format and compressed. The columnar storage format enables efficient querying and data retrieval, as it allows for reading only the necessary columns for a query, thereby reducing IO operations. Additionally, Snowflake uses advanced compression techniques to minimize storage costs and improve performance. This combination of columnar storage and compression is key to Snowflake ' s data warehousing capabilities.
References:
Snowflake Documentation: Data Storage and Organization

**NEW QUESTION: 516**
How does clustering depth impact query performance?
**A.** Low clustering depth indicates that there are many overlapping micro-partitions and that query pruning is efficient.

**B.** High clustering depth indicates that queries are taking advantage of caching and are running efficiently.
**C.** Low clustering depth indicates that there are few overlapping micro-partitions and that query pruning is efficient.
**D.** High clustering depth indicates that too many users are scanning the same micro-partitions, slowing query performance.
**Answer: (SHOW ANSWER)**
Clustering depth measures the average number of overlapping micro-partitions for specified columns.
A smaller clustering depth means fewer micro-partitions overlap, indicating that related values are stored in a more ordered and localized manner. This enables Snowflake to prune more irrelevant micro- partitions when a query filters on those columns, reducing the amount of data scanned and generally improving performance. Therefore, option C correctly states that low clustering depth corresponds to relatively little overlap and efficient pruning. High clustering depth indicates greater overlap and can cause more micro-partitions to be scanned. Snowflake notes that clustering depth is a diagnostic measurement rather than an absolute performance guarantee. Actual query performance remains the final indicator of whether a table is sufficiently well clustered.

**NEW QUESTION: 517**
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: 518**
Which table function will automatically detect the file metadata schema for a set of staged semi-structured data files, and then retrieve the column definitions?
**A.** GET_OBJECT_REFERENCES
**B.** SPLIT_TO_TABLE
**C.** VALIDATE
**D.** INFER_SCHEMA
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 519**
Which statements describe benefits of Snowflake ' s separation of compute and storage? (Select TWO).
**A.** The separation allows independent scaling of computing resources.
**B.** The separation ensures consistent data encryption across all virtual data warehouses.
**C.** The separation supports automatic conversion of semi-structured data into structured data for advanced data analysis.
**D.** Storage volume growth and compute usage growth can be tightly coupled.

**E.** Compute can be scaled up or down without the requirement to add more storage.

**Answer: (SHOW ANSWER)**

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

**NEW QUESTION: 520**

What is the only supported character set for loading and unloading data from all supported file formats?

**A.** UTF-8

**B.** UTF-16

**C.** ISO-8859-1

**D.** WINDOWS-1253

**Answer: (SHOW ANSWER)**

UTF-8 is the only supported character set for loading and unloading data from all supported file formats in Snowflake. UTF-8 is a widely used encoding that supports a large range of characters from various languages, making it suitable for internationalization and ensuring data compatibility across different systems and platforms.

References:

Snowflake Documentation: Data Loading and Unloading

**NEW QUESTION: 521**

If a column contains float values that are greater than (15,9) and the values need to be preserved when unloading data, what data file format should be used?

**A.** JSON

**B.** PARQUET

**C.** AVRO

**D.** ORC

**Answer: A (LEAVE A REPLY)**

The correct answer is A. JSON .

When unloading floating-point or high-precision numeric values, users must consider whether the target file format can preserve the values as needed. JSON stores values in a textual semi-structured representation and can preserve numeric values without forcing them into a fixed Parquet numeric precision representation.

Why A is correct:

For values that exceed a precision/scale such as (15,9) and must be preserved during unloading, JSON is the safest option among the listed choices because it can represent the values textually in the output.

Why the other options are incorrect:

B). PARQUET has limitations for certain numeric precision/scale representations during unload. Snowflake documentation notes precision-related considerations for Parquet unloads.

C). AVRO is a semi-structured file format, but it is not the best documented answer for preserving float values beyond this precision scenario.

D). ORC is not the best answer for this Snowflake unload precision scenario.

Official Snowflake documentation reference:

Snowflake documentation for unloading data explains that file format choice can affect how numeric values are represented and preserved. JSON is commonly used when preserving semi-structured or text-based representations is required.
Reference: Snowflake Documentation - Unloading data; Snowflake Documentation - File format considerations; SnowPro Core Study Guide - Data Loading and Unloading.
==

**NEW QUESTION: 522**
Which Snowflake tool or feature enables governed sharing of datasets with a selected group of consumers that are invited, internal or external, to a Snowflake user ' s company?
**A.** Replication
**B.** Snowpipe
**C.** Data Exchange
**D.** Snowflake Marketplace
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 523**
Which Snowflake role can manage any object grant globally, including modifying and revoking grants?
**A.** USERADMIN
**B.** ORGADMIN
**C.** SYSADMIN
**D.** SECURITYADMIN
**Answer: (SHOW ANSWER)**
The SECURITYADMIN role in Snowflake can manage any object grant globally, including modifying and revoking grants. This role has the necessary privileges to oversee and control access to all securable objects within the Snowflake environment4.

**NEW QUESTION: 524**
Which Snowflake layer is always leveraged when accessing a query from the result cache?
**A.** Metadata
**B.** Data Storage
**C.** Compute
**D.** Cloud Services
**Answer: (SHOW ANSWER)**
The Cloud Services layer in Snowflake is responsible for managing the result cache.When a query is executed, the results are stored in this cache, and subsequent identical queries can leverage these cached results without re-executing the entire query1.

**NEW QUESTION: 525**
What is a feature of column-level security in Snowflake?
**A.** Role access policies

**B.** Network policies
**C.** Internal tokenization
**D.** External tokenization
**Answer: A (LEAVE A REPLY)**
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: 526**
When an ACCOUNTADMIN gives a user a custom role, what privilege or privileges is the user granted by default?
**A.** All privileges on objects allowed by the custom role
**B.** Access to all the objects owned by the USERADMIN role
**C.** Access to the PUBLIC role
**D.** All privileges that have been granted to the ACCOUNTADMIN
**Answer: C (LEAVE A REPLY)**

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

**NEW QUESTION: 527**
Which type of loop requires a BREAK statement to stop executing?
**A.** FOR
**B.** LOOP
**C.** REPEAT
**D.** WHILE
**Answer: (SHOW ANSWER)**
The LOOP type of loop in Snowflake Scripting does not have a built-in termination condition and requires a BREAK statement to stop executing4.

**NEW QUESTION: 528**
Which Snowflake function is maintained separately from the data and helps to support features such as Time Travel, Secure Data Sharing, and pruning?

**A.** Column compression
**B.** Data clustering
**C.** Micro-partitioning
**D.** Metadata management
**Answer: (SHOW ANSWER)**
Micro-partitioning is a Snowflake function that is maintained separately from the data and supports features such as Time Travel, Secure Data Sharing, and pruning.It allows Snowflake to efficiently manage and query large datasets by organizing them into micro-partitions1.

**NEW QUESTION: 529**
In the query profiler view for a query, which components represent areas that can be used to help optimize query performance? (Select TWO)
**A.** Bytes scanned
**B.** Bytes sent over the network
**C.** Number of partitions scanned
**D.** Percentage scanned from cache
**E.** External bytes scanned
**Answer: (SHOW ANSWER)**
In the query profiler view, the components that represent areas that can be used to help optimize query performance include 'Bytes scanned' and 'Number of partitions scanned'. 'Bytes scanned' indicates the total amount of data the query had to read and is a direct indicator of the query's efficiency. Reducing the bytes scanned can lead to lower data transfer costs and faster query execution. 'Number of partitions scanned' reflects how well the data is clustered; fewer partitions scanned typically means better performance because the system can skip irrelevant data more effectively.
References:
[COF-C03] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Query Profiling1

**NEW QUESTION: 530**
Which statements are correct concerning the leveraging of third-party data from the Snowflake Data Marketplace? (Choose two.)
Freecram.net
**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 movingit to an individual Snowflake account, allowing for seamless integration with existing data

**NEW QUESTION: 531**

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: 532**
What do the terms scale up and scale out refer to in Snowflake? (Choose two.)
**A.** Scaling out adds clusters of the same size to a virtual warehouse to handle more concurrent queries.
**B.** Scaling out adds clusters of varying sizes to a virtual warehouse.
**C.** Scaling out adds additional database servers to an existing running cluster to handle more concurrent queries.
**D.** Snowflake recommends using both scaling up and scaling out to handle more concurrent queries.
**E.** Scaling up resizes a virtual warehouse so it can handle more complex workloads.
**F.** Scaling up adds additional database servers to an existing running cluster to handle larger workloads.
**Answer: (SHOW ANSWER)**
Scaling out in Snowflake involves adding clusters of the same size to a virtual warehouse, which allows for handling more concurrent queries without affecting the performance of individual queries. Scaling up refers to resizing a virtual warehouse to increase its compute resources, enabling it to handle more complex workloads and larger queries more efficiently.

**NEW QUESTION: 533**
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: 534**
Which virtual warehouse parameters are used to configure how clusters run in Maximized mode or Auto-scale mode? (Choose two.)

**A.** WAREHOUSE_SIZE
**B.** MIN_CLUSTER_COUNT
**C.** MAX_CONCURRENCY_LEVEL
**D.** WAREHOUSE_TYPE
**E.** MAX_CLUSTER_COUNT
**Answer: (SHOW ANSWER)**
MIN_CLUSTER_COUNT and MAX_CLUSTER_COUNT determine how many clusters a multi- cluster virtual warehouse can run. When the two values are equal and greater than one, the warehouse operates in Maximized mode and starts all configured clusters. When MIN_CLUSTER_COUNT is lower than MAX_CLUSTER_COUNT, the warehouse operates in Auto-scale mode and starts or stops clusters within that range as the workload changes. WAREHOUSE_SIZE controls the compute capacity of each cluster rather than the operating mode. MAX_CONCURRENCY_LEVEL limits the number of concurrently running statements per cluster, and WAREHOUSE_TYPE identifies the warehouse type rather than its scaling boundaries. Therefore, the two parameters required to configure Maximized or Auto-scale behavior are MIN_CLUSTER_COUNT and MAX_CLUSTER_COUNT.

**NEW QUESTION: 535**
Which option can be used to define column data types when unloading data into Parquet-formatted files?
**A.** BINARY_AS_TEXT = TRUE
**B.** EMPTY_FIELD_AS_NULL = FALSE
**C.** CAST or the :: operator in the COPY INTO statement
**D.** SINGLE = TRUE
**Answer: (SHOW ANSWER)**
COPY INTO < location > can unload the result of a SELECT query rather than exporting a table directly. Expressions in that SELECT query can use CAST or the :: operator to convert columns to the data types required in the generated Parquet file. For example, a numeric expression can be cast to a specific NUMBER type before it is written. BINARY_AS_TEXT controls how binary values are represented, EMPTY_FIELD_AS_NULL affects the treatment of empty fields in applicable formats, and SINGLE controls whether Snowflake generates one output file or multiple files. None of those options generally defines the data type of an output column. Snowflake advises confirming that every source value fits the target type because casting to an insufficient size can truncate values.

**NEW QUESTION: 536**
Which command should be used to drop files from an internal or external stage?
**A.** DROP
**B.** DELETE
**C.** REMOVE
**D.** TRUNCATE
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 537**
What command will list all of the roles that have been granted to role DATA_ENGINEER?
**A.** SHOW GRANTS OF ROLE DATA_ENGINEER;

**B.** SHOW GRANTS ON ROLE DATA_ENGINEER;
**C.** SHOW GRANTS TO ROLE DATA_ENGINEER;
**D.** SHOW ROLES LIKE ' DATA_ENGINEER ' ;
**Answer: (SHOW ANSWER)**

**NEW QUESTION: 538**
How should clustering be used to optimize the performance of queries that run on a very large table?
**A.** Manually re-cluster the table regularly.
**B.** Choose one high cardinality column as the clustering key.
**C.** Use the column that is most-frequently used in query select clauses as the clustering key.
**D.** Assess the average table depth to identify how clustering is impacting the query.
**Answer: (SHOW ANSWER)**
For optimizing the performance of queries that run on a very large table, it is recommended to choose one high cardinality column as the clustering key. This helps to co-locate similar rows in thesame micro- partitions, improving scan efficiency in queries by skipping data that does not match filtering predicates4.
References: [ COF-C03 ] SnowPro Core Certification Exam Study Guide

**NEW QUESTION: 539**
By default, which system roles can create and manage users? (Select TWO)
**A.** USERADMIN
**B.** SYSADMIN
**C.** SECURITYADMIN
**D.** PUBLIC
**E.** ORGADMIN
**Answer: (SHOW ANSWER)**
* USERADMIN : Used for managing users and roles.
* SECURITYADMIN : Can manage users, roles, and access control.
SYSADMIN is used primarily for object creation, not user management.
Reference:
Snowflake Docs: System-Defined Roles

**NEW QUESTION: 540**
Which Snowflake feature, interface, or service is used to create and share custom web applications for machine learning and data science?
Freecram.net
**A.** Cortex AI
**B.** Streamlit in Snowflake
**C.** Snowpipe
**D.** Snowsight
**Answer: (SHOW ANSWER)**
Streamlit in Snowflake is the feature intended for building and sharing interactive web applications for machine learning, analytics, and data science. Streamlit is a Python application framework that allows developers to turn scripts and data

workflows into browser-based applications with interactive controls and visualizations. Running Streamlit within Snowflake keeps governed data and much of the application processing inside the Snowflake environment. Cortex AI supplies AI models and functions, but it is not the web-application framework requested. Snowpipe is a continuous data-ingestion service, and Snowsight is Snowflake's broader web interface for worksheets, administration, monitoring, and development. Although Streamlit applications are created and managed through Snowflake interfaces, the specific application-building service is Streamlit in Snowflake. Therefore, option B is correct.

**NEW QUESTION: 541**
Snowflake's hierarchical key mode includes which keys? (Select TWO).
**A.** Account master keys
**B.** Database master keys
**C.** File keys
**D.** Secure view keys
**E.** Schema master keys
**Answer: (SHOW ANSWER)**
Snowflake's hierarchical key model includes several levels of keys, where Account master keys and File keys are part of this hierarchy.Account master keys are used to encrypt all the data within an account, while File keys are used to encrypt individual files within the database2.

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

**NEW QUESTION: 542**
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: (SHOW ANSWER)**
https://docs.snowflake.com/en/user-guide/admin-security-privatelink.html

**NEW QUESTION: 543**
Which Snowflake feature supports the use of natural language queries to analyze unstructured data?
**A.** Snowsight
**B.** Snowflake ML

**C.** Snowflake Cortex
**D.** Snowflake Horizon Catalog
**Answer: (SHOW ANSWER)**
The correct answer is C. Snowflake Cortex .
Snowflake Cortex provides AI and large language model capabilities inside Snowflake. It includes features that support natural language interaction, text analysis, document processing, search, summarization, classification, and other AI-powered workflows.
Why C is correct:
Snowflake Cortex is the Snowflake AI feature family used for natural language and unstructured data use cases. Cortex capabilities can be used to analyze text and other unstructured data using natural language functions and AI services.
Why the other options are incorrect:
A). Snowsight is Snowflake's web interface. It can expose AI-assisted features, but the underlying AI capability is Snowflake Cortex.
B). Snowflake ML supports machine learning workflows but is not the primary feature family for natural language querying of unstructured data.
D). Snowflake Horizon Catalog is focused on governance, discovery, and catalog capabilities, not natural language analysis of unstructured data.
Official Snowflake documentation reference:
Snowflake documentation describes Snowflake Cortex as the set of intelligent, AI-powered features that use large language models and ML to analyze data, including unstructured text data.
Reference: Snowflake Documentation - Snowflake Cortex; Snowflake Documentation - Cortex AI functions; SnowPro Core Study Guide - Snowflake Architecture.
==

**NEW QUESTION: 544**
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: (SHOW ANSWER)**
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: 545**
Which command is used when loading data with Snowpipe?

**A.** INSERT
**B.** COPY
**C.** UPDATE
**D.** MERGE
**Answer: (SHOW ANSWER)**
The correct answer is B. COPY .
Snowpipe loads data using a pipe object. The pipe definition contains a COPY INTO < table > statement that tells Snowflake how to load staged files into a target table.
Why B is correct:
Snowpipe automates loading staged files, but the actual load operation is still based on the COPY INTO < table > command.
Example:
CREATE PIPE my_pipe
AS
COPY INTO my_table
FROM @my_stage
FILE_FORMAT = (TYPE = CSV);
Why the other options are incorrect:
A). INSERT adds rows manually or through query results, but Snowpipe does not use INSERT as its load command.
C). UPDATE modifies existing rows and is not used by Snowpipe for loading files.
D). MERGE performs insert/update/delete logic based on matching conditions, but Snowpipe uses COPY INTO < table > for file ingestion.
Official Snowflake documentation reference:
Snowflake documentation explains that a pipe defines a COPY INTO < table > statement used by Snowpipe to load data from staged files.
Reference: Snowflake Documentation - Snowpipe; Snowflake Documentation - CREATE PIPE; Snowflake Documentation - COPY INTO < table > ; SnowPro Core Study Guide - Data Loading and Unloading.
==

**NEW QUESTION: 546**
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 connectto 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: 547**
A stream can be created on which Snowflake objects to record data manipulation language(DML) changes? (Select TWO).
**A.** Database
**B.** Standard tables
**C.** Standard tables
**D.** Standard views
**E.** Schemas
**F.** Pipes
**Answer: (SHOW ANSWER)**
Snowflake streams are objects that enable users to track changes (inserts, updates, and deletes) to the data in tables, facilitating real-time data processing and integration workflows. Streams can be created on standard tables, capturing DML changes made to these tables so that downstream processes can consume the changes incrementally. This feature supports efficient data ETL, replication, and real-time analytics by providing a mechanism to process only the data that has changed. Note: The correct options should likely include a distinction between " Standard tables " and another object type such as " External tables " rather than repeating " Standard tables " twice.
References: Snowflake Documentation on Streams

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