Snowflake

Exam Questions COF-C02

SnowPro Core Certification Exam (COF-C02)

NEW QUESTION 1
- (Topic 1)
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: A

Explanation:
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 2
- (Topic 1)
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: C

Explanation:
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-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Snowpipe1

NEW QUESTION 3
- (Topic 1)
What is the default character set used when loading CSV files into Snowflake?

A. UTF-8
B. UTF-16
C. ISO S859-1
D. ANSI_X3.A

Answer: A

Explanation:
https://docs.snowflake.com/en/user-guide/intro-summary-loading.html#:~:text=For%20delimited%20files%20(CSV%2C%20TSV,encoding%20to%20use%20for%20loading.
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 4
- (Topic 1)
What can be used to view warehouse usage over time? (Select Two).

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

Answer: BD

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

NEW QUESTION 5
- (Topic 1)
True or False: Reader Accounts are able to extract data from shared data objects for use outside of Snowflake.

A. True
B. False

Answer: B

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

NEW QUESTION 6
- (Topic 1)
What Snowflake features allow virtual warehouses to handle high concurrency workloads? (Select TWO)

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

Answer: BD

Explanation:
Snowflake??s architecture is designed to handle high concurrency workloads through several features, two of which are particularly effective:
? B. The use of warehouse auto scaling: This feature allows Snowflake to automatically adjust the compute resources allocated to a virtual warehouse in response to the workload. If there is an increase in concurrent queries, Snowflake can scale up the resources to maintain performance.
? D. Use of multi-clustered warehouses: Multi-clustered warehouses enable Snowflake to run multiple clusters of compute resources simultaneously. This allows for the distribution of queries across clusters, thereby reducing the load on any single cluster and improving the system??s ability to handle a high number of concurrent queries.
These features ensure that Snowflake can manage varying levels of demand without manual intervention, providing a seamless experience even during peak usage. References:
? Snowflake Documentation on Virtual Warehouses
? SnowPro® Core Certification Study Guide

NEW QUESTION 7
- (Topic 1)
A company strongly encourages all Snowflake users to self-enroll in Snowflake's default Multi-Factor Authentication (MFA) service to provide increased login security for users connecting to Snowflake.
Which application will the Snowflake users need to install on their devices in order to connect with MFA?

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

Answer: B

Explanation:
Snowflake??s default Multi-Factor Authentication (MFA) service is powered by Duo Security. Users are required to install the Duo Mobile application on their devices to
use MFA for increased login security when connecting to Snowflake. This service is managed entirely by Snowflake, and users do not need to sign up separately with Duo1.

NEW QUESTION 8
- (Topic 1)
What are value types that a VARIANT column can store? (Select TWO)

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

Answer: BD

Explanation:
A VARIANT column in Snowflake can store semi-structured data types. This includes:
? B. OBJECT: An object is a collection of key-value pairs in JSON, and a VARIANT column can store this type of data structure.
? D. ARRAY: An array is an ordered list of zero or more values, which can be of any variant-supported data type, including objects or other arrays.
The VARIANT data type is specifically designed to handle semi-structured data like JSON, Avro, ORC, Parquet, or XML, allowing for the storage of nested and complex data structures.
References:
? Snowflake Documentation on Semi-Structured Data Types
? SnowPro® Core Certification Study Guide

NEW QUESTION 9

- (Topic 1)
What is the recommended file sizing for data loading using Snowpipe?

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

Answer: C

Explanation:
For data loading using Snowpipe, the recommended file size is a compressed file greater than 10 MB and up to 100 MB. This size range is optimal for Snowpipe??s continuous, micro-batch loadingprocess, allowing for efficient and timely data ingestion without overwhelming the system with files that are too large or too small. References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Snowpipe1

NEW QUESTION 10
- (Topic 1)
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: B

Explanation:
To find information about query activity from 90 days ago, a Snowflake user should use the account_usage.query_history_archive view. This view is designed to provide access to historical query data beyond the default 14-day retention period found in the standard query_history view. It allows users to analyze and audit past query activities for up to 365 days after the date of execution, which includes the 90-day period mentioned. References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Account Usage Schema1

NEW QUESTION 10
- (Topic 1)
The fail-safe retention period is how many days?

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

Answer: B

Explanation:
Fail-safe is a feature in Snowflake that provides an additional layer of data protection. After the Time Travel retention period ends, Fail-safe offers a non-configurable 7-day period during which historical data may be recoverable by Snowflake. This period is designed to protect against accidental data loss and is not intended for customer access. References: Understanding and viewing Fail-safe | Snowflake Documentation

NEW QUESTION 13
- (Topic 1)
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

Explanation:
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 18
- (Topic 1)
What is the purpose of an External Function?

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

Answer: A

Explanation:
The purpose of an External Function in Snowflake is to call code that executes outside of the Snowflake environment. This allows Snowflake to interact with external services and leverage functionalities that are not natively available within Snowflake, such as calling APIs or running custom code hosted on cloud services3. https://docs.snowflake.com/en/sql-reference/external-functions.html

NEW QUESTION 20
- (Topic 1)
Which Snowflake object enables loading data from files as soon as they are available in a cloud storage location?

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

Answer: A

Explanation:
In Snowflake, a Pipe is the object designed to enable the continuous, near- real-time loading of data from files as soon as they are available in a cloud storage location. Pipes use Snowflake??s COPY command to load data and can be associated with a Stage object to monitor for new files. When new data files appear in the stage, the pipe automatically loads the data into the target table.
References:
? Snowflake Documentation on Pipes
? SnowPro® Core Certification Study Guide https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro.html

NEW QUESTION 25
- (Topic 1)
Which of the following compute resources or features are managed by Snowflake? (Select TWO).

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

Answer: CE

Explanation:
Snowflake manages various compute resources and features, including Snowpipe and the ability to scale up a warehouse. Snowpipe is Snowflake??s continuous data ingestion service that allows users to load data as soon as it becomes available. Scaling up a warehouse refers to increasing the compute resources allocated to a virtual warehouse to handle larger workloads or improve performance.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Snowpipe and Virtual Warehouses1

NEW QUESTION 26
- (Topic 1)
What is the minimum Snowflake edition required to create a materialized view?

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

Answer: B

Explanation:
Materialized views in Snowflake are a feature that allows for the pre- computation and storage of query results for faster query performance. This feature is available starting from the Enterprise Edition of Snowflake. It is not available in the Standard Edition, and while it is also available in higher editions like Business Critical and Virtual Private Snowflake, the Enterprise Edition is the minimum requirement. References:
? Snowflake Documentation on CREATE MATERIALIZED VIEW1.
? Snowflake Documentation on Working with Materialized Views https://docs.snowflake.com/en/sql-reference/sql/create-materialized-view.html#:~:text=Materialized%20views%20require%20Enterprise%20Edition,upgrading%2C%20please%20contact%20Snowflake%20Support.

NEW QUESTION 31
- (Topic 1)
What is a key feature of Snowflake architecture?

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

Answer: C

Explanation:

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

NEW QUESTION 33
- (Topic 2)
What features that are part of the Continuous Data Protection (CDP) feature set in Snowflake do not require additional configuration? (Choose two.)

A. Row level access policies
B. Data masking policies
C. Data encryption
D. Time Travel
E. External tokenization

Answer: CD

Explanation:
Data encryption and Time Travel are part of Snowflake??s Continuous Data Protection (CDP) feature set that do not require additional configuration. Data encryption is automatically applied to all filesstored on internal stages, and Time Travel allows for querying and restoring data without any extra setup

NEW QUESTION 37
- (Topic 2)
Which tasks are performed in the Snowflake Cloud Services layer? (Choose two.)

A. Management of metadata
B. Computing the data
C. Maintaining Availability Zones
D. Infrastructure security
E. Parsing and optimizing queries

Answer: AE

Explanation:
The Snowflake Cloud Services layer performs a variety of tasks, including the management of metadata and the parsing and optimization of queries. This layer is responsible for coordinating activities across Snowflake, including user session management, security, and query compilation3.

NEW QUESTION 40
- (Topic 2)
What affects whether the query results cache can be used?

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

Answer: C

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

NEW QUESTION 41
- (Topic 2)
In a Snowflake role hierarchy, what is the top-level role?

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

Answer: C

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

NEW QUESTION 42
- (Topic 2)
Which command sets the Virtual Warehouse for a session?

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

Answer: C

Explanation:
The command USE WAREHOUSE <<warehouse name>>; is used to set the virtual warehouse for the current session in Snowflake. This command specifies which virtual warehouse to use for executing queries in that session1.
Reference: https://docs.snowflake.com/en/user-guide/warehouses-tasks.html

NEW QUESTION 44
- (Topic 2)
When cloning a database containing stored procedures and regular views, that have fully qualified table references, which of the following will occur?

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

Answer: A

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

NEW QUESTION 46
- (Topic 2)
A company needs to allow some users to see Personally Identifiable Information (PII) while limiting other users from seeing the full value of the PII.
Which Snowflake feature will support this?

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

Answer: B

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

NEW QUESTION 51
- (Topic 2)
The Snowflake Search Optimization Services supports improved performance of which kind of query?

A. Queries against large tables where frequent DML occurs
B. Queries against tables larger than 1 TB
C. Selective point lookup queries
D. Queries against a subset of columns in a table

Answer: C

Explanation:
The Snowflake Search Optimization Service is designed to support improved performance for selective point lookup queries. These are queries that retrieve specific records from a database, often based on a unique identifier or a small set of criteria3.

NEW QUESTION 55
- (Topic 2)
Which of the following are characteristics of Snowflake virtual warehouses? (Choose two.)

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

Answer: CE

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

NEW QUESTION 57
- (Topic 2)
Which Snowflake architectural layer is responsible for a query execution plan?

A. Compute
B. Data storage

C. Cloud services
D. Cloud provider

Answer: C

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

NEW QUESTION 60
- (Topic 2)
If 3 size Small virtual warehouse is made up of two servers, how many servers make up a
Large warehouse?

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

Answer: B

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

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

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

NEW QUESTION 64
- (Topic 2)
Network policies can be set at which Snowflake levels? (Choose two.)

A. Role
B. Schema
C. User
D. Database
E. Account
F. Tables

Answer: CE

Explanation:
Network policies in Snowflake can be set at the user level and at the account level2.
Reference: https://docs.snowflake.com/en/user-guide/network-policies.html#creating-network-policies

NEW QUESTION 66
- (Topic 2)
True or False: Snowpipe via REST API can only reference External Stages as source.

A. True
B. False

Answer: B

Explanation:
Snowpipe via REST API can reference both named internal stages within Snowflake and external stages, such as Amazon S3, Google Cloud Storage, or Microsoft Azure1. This means that Snowpipe is not limited to only external stages as a source for data loading.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation1
Reference: https://community.snowflake.com/s/article/Making-Transient-table-by-Default

NEW QUESTION 68
- (Topic 2)
Which statement is true about running tasks in Snowflake?

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

Answer: B

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

NEW QUESTION 70
- (Topic 2)
What are supported file formats for unloading data from Snowflake? (Choose three.)

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

Answer: BCF

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

NEW QUESTION 73
- (Topic 2)
What is the minimum Snowflake edition that has column-level security enabled?

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

Answer: B

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

NEW QUESTION 78
- (Topic 2)
In an auto-scaling multi-cluster virtual warehouse with the setting SCALING_POLICY = ECONOMY enabled, when is another cluster started?

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

Answer: A

Explanation:
In an auto-scaling multi-cluster virtual warehouse with the SCALING_POLICY set to ECONOMY, another cluster is started when the system has enough load for 2 minutes (A). This policy is designed to optimize the balance between performance and cost, starting additional clusters only when the sustained load justifies it2.

NEW QUESTION 80
- (Topic 2)
Which of the following accurately describes shares?

A. Tables, secure views, and secure UDFs can be shared
B. Shares can be shared
C. Data consumers can clone a new table from a share
D. Access to a share cannot be revoked once granted

Answer: A

Explanation:
Shares in Snowflake are named objects that encapsulate all the information required to share databases, schemas, tables, secure views, and secure UDFs. These objects can be added to a share by granting privileges on them to the share via a database role

NEW QUESTION 81
- (Topic 2)
A virtual warehouse is created using the following command:
Create warehouse my_WH with warehouse_size = MEDIUM min_cluster_count = 1
max_cluster_count = 1
auto_suspend = 60 auto_resume = true;
The image below is a graphical representation of the warehouse utilization across two days.

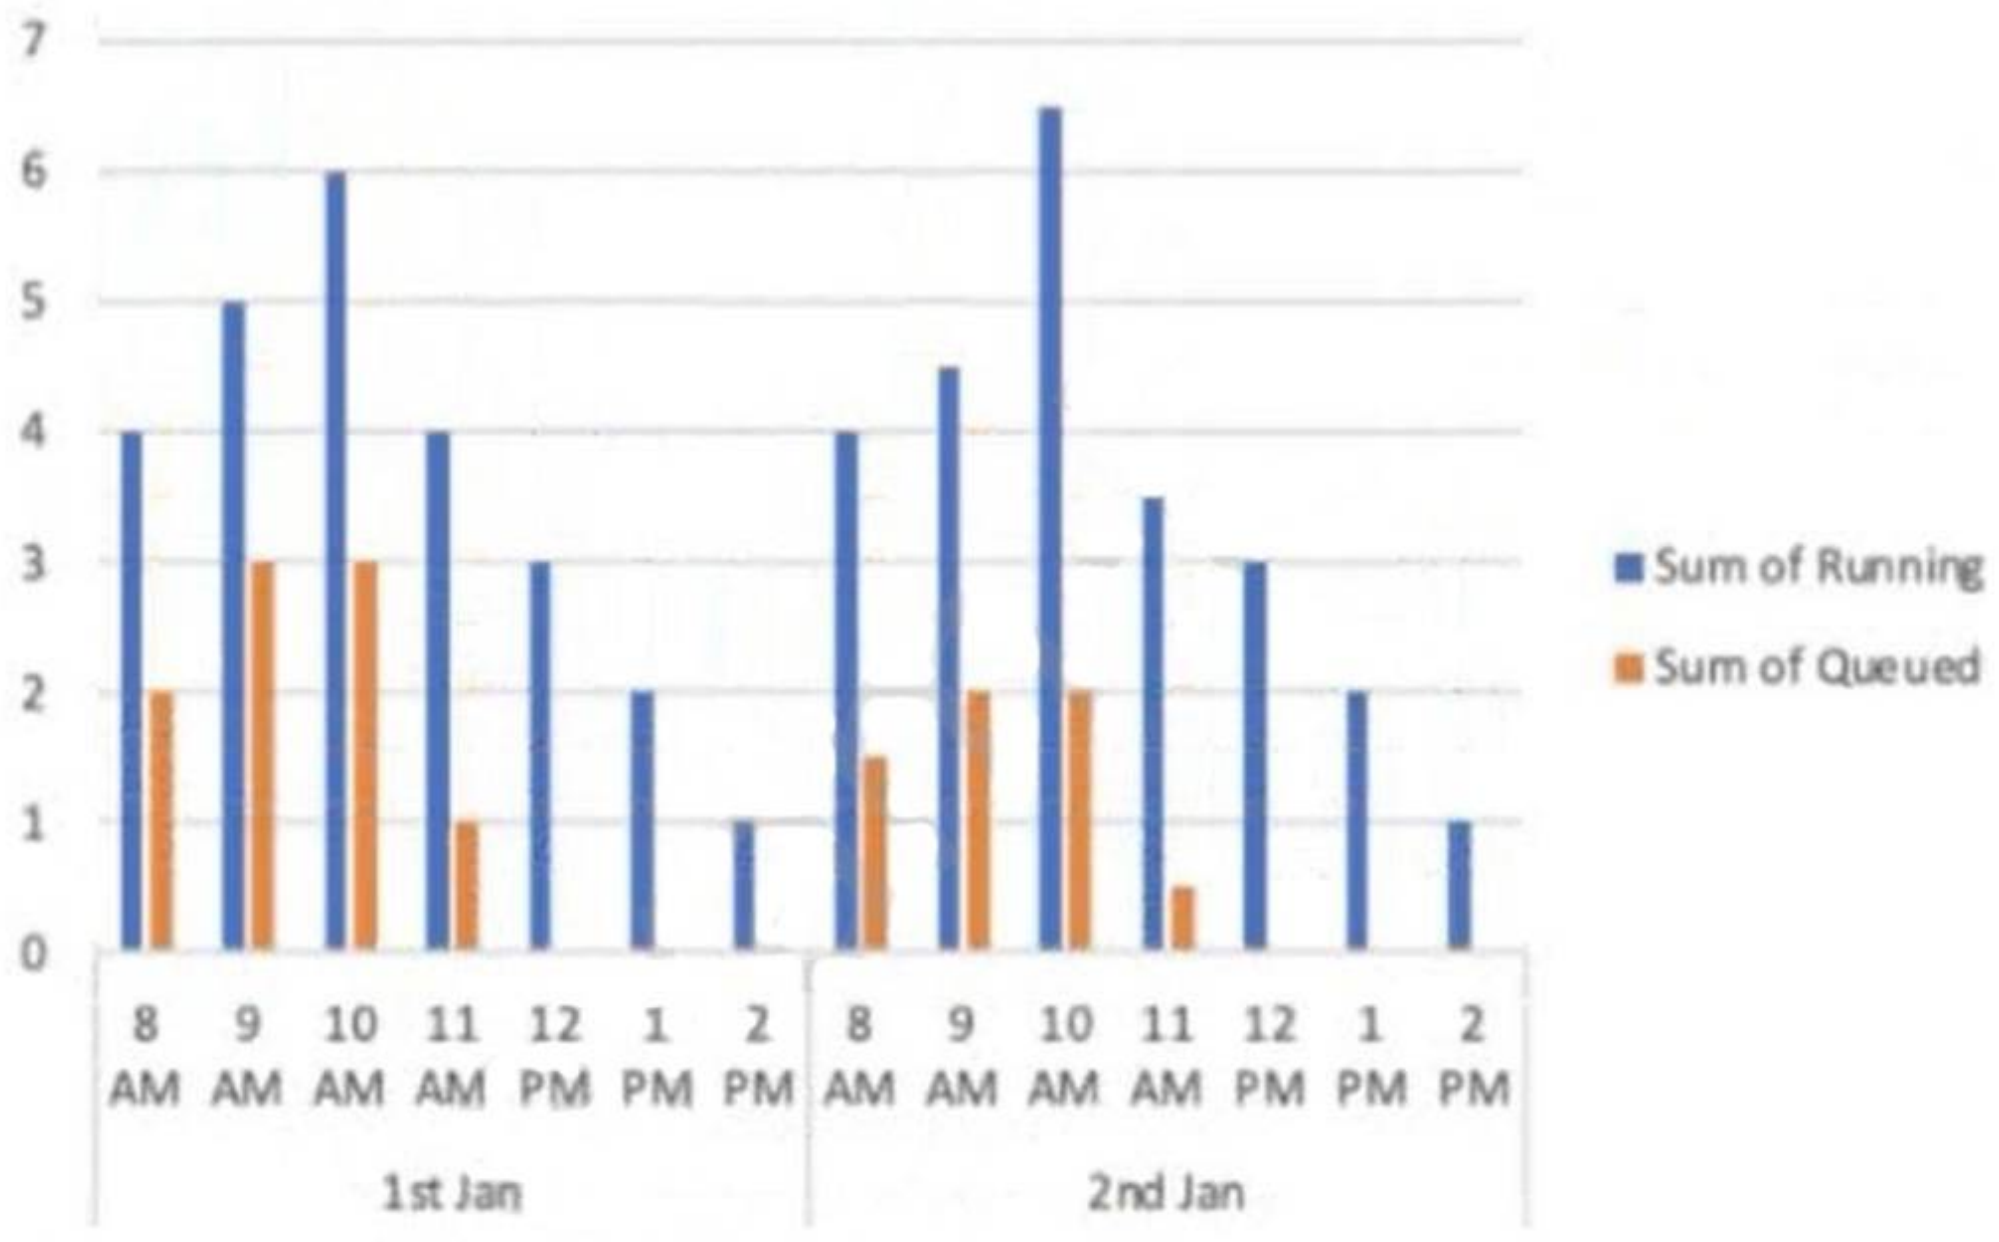

What action should be taken to address this situation?

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

Answer: C

Explanation:
The graphical representation of warehouse utilization indicates periods of significant queuing, suggesting that the current single cluster cannot efficiently handle all incoming queries. Configuring the warehouse to a multi-cluster warehouse will distribute the load among multiple clusters, reducing queuing times and improving overall performance1.
References = Snowflake Documentation on Multi-cluster Warehouses1

NEW QUESTION 83
- (Topic 2)
What is the following SQL command used for? Select * from table(validate(t1, job_id => '_last'));

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

Answer: D

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

NEW QUESTION 84
- (Topic 2)
Why does Snowflake recommend file sizes of 100-250 MB compressed when loading data?

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

Answer: D

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

NEW QUESTION 89
- (Topic 2)
A user has unloaded data from a Snowflake table to an external stage.
Which command can be used to verify if data has been uploaded to the external stage named my_stage?

A. view @my_stage
B. list @my_stage
C. show @my_stage
D. display @my_stage

Answer: B

Explanation:
The list @my_stage command in Snowflake can be used to verify if data has been uploaded to an external stage named my_stage. This command provides a list of files that are present in the specified stage2.

NEW QUESTION 94
- (Topic 3)
Which parameter prevents streams on tables from becoming stale?

A. MAXDATAEXTENSIONTIMEINDAYS
B. MTN_DATA_RETENTION_TTME_TN_DAYS
C. LOCK_TIMEOUT
D. STALE_AFTER

Answer: A

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

NEW QUESTION 99
- (Topic 3)
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: C

Explanation:
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 104
- (Topic 3)
Which user object property requires contacting Snowflake Support in order to set a value for it?

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

Answer: B

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

NEW QUESTION 106
- (Topic 3)
If a multi-cluster warehouse is using an economy scaling policy, how long will queries wait in the queue before another cluster is started?

A. 1 minute
B. 2 minutes
C. 6 minutes
D. 8 minutes

Answer: B

Explanation:
In a multi-cluster warehouse with an economy scaling policy, queries will wait in the queue for 2 minutes before another cluster is started. This is to minimize costs by allowing queries to queue up for a short period before adding additional compute resources. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 107
- (Topic 3)
Which command is used to unload files from an internal or external stage to a local file system?

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

Answer: B

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

NEW QUESTION 112
- (Topic 3)
How can a user change which columns are referenced in a view?

A. Modify the columns in the underlying table
B. Use the ALTER VIEW command to update the view
C. Recreate the view with the required changes
D. Materialize the view to perform the changes

Answer: C

Explanation:
In Snowflake, to change the columns referenced in a view, the view must be recreated with the required changes. The ALTER VIEW command does not allow changing the definition of a view; it can only be used to rename a view, convert it to or from a secure view, or add, overwrite, or remove a comment for a view. Therefore, the correct approach is to drop the existing view and create a new one with the desired column references.

NEW QUESTION 114
- (Topic 3)
Which formats does Snowflake store unstructured data in? (Choose two.)

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

Answer: AC

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

NEW QUESTION 117
- (Topic 3)
What service is provided as an integrated Snowflake feature to enhance Multi-Factor Authentication (MFA) support?

A. Duo Security
B. OAuth
C. Okta
D. Single Sign-On (SSO)

Answer: A

Explanation:
Snowflake provides Multi-Factor Authentication (MFA) support as an integrated feature, powered by the Duo Security service. This service is managed completely by Snowflake, and users do not need to sign up separately with Duo1

NEW QUESTION 121
- (Topic 3)
What is a characteristic of the Snowflake Query Profile?

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

Answer: B

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

NEW QUESTION 123
- (Topic 3)
A materialized view should be created when which of the following occurs? (Choose two.)

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

Answer: BE

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

NEW QUESTION 128
- (Topic 3)
Which transformation is supported by a COPY INTO <table> command?

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

Answer: C

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

NEW QUESTION 129
- (Topic 3)
Which stages are used with the Snowflake PUT command to upload files from a local file system? (Choose three.)

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

Answer: BDF

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

NEW QUESTION 133
- (Topic 3)
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: B

Explanation:

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 137
- (Topic 3)
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: BE

Explanation:
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 141
- (Topic 3)
What is the recommended compressed file size range for continuous data loads using Snowpipe?

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

Answer: D

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

NEW QUESTION 145
- (Topic 3)
If file format options are specified in multiple locations, the load operation selects which option FIRST to apply in order of precedence?

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

Answer: D

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

NEW QUESTION 147
- (Topic 3)
Which statement describes pruning?

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

Answer: A

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

NEW QUESTION 148
- (Topic 3)
Which of the following statements describes a schema in Snowflake?

A. A logical grouping of objects that belongs to a single database
B. A logical grouping of objects that belongs to multiple databases
C. A named Snowflake object that includes all the information required to share a database
D. A uniquely identified Snowflake account within a business entity

Answer: A

Explanation:
A schema in Snowflake is a logical grouping of database objects, such as tables and views, that belongs to a single database. Each schema is part of a

namespace in Snowflake, which is inferred from the current database and schema in use for the session5

NEW QUESTION 149
- (Topic 3)
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: B

Explanation:
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 151
- (Topic 3)
How long can a data consumer who has a pre-signed URL access data files using Snowflake?

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

Answer: D

Explanation:
A data consumer who has a pre-signed URL can access data files using Snowflake until the expiration time is exceeded. The expiration time is set when the pre-signed URL is generated and determines how long the URL remains valid3.

NEW QUESTION 152
- (Topic 3)
Which kind of Snowflake table stores file-level metadata for each file in a stage?

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

Answer: A

Explanation:
The kind of Snowflake table that stores file-level metadata for each file in a stage is a directory table. A directory table is an implicit object layered on a stage and stores file-level metadata about the data files in the stage3.

NEW QUESTION 153
- (Topic 3)
When unloading data to an external stage, what is the MAXIMUM file size supported?

A. 1 GB
B. 5 GB
C. 10 GB
D. 16 GB

Answer: B

Explanation:
When unloading data to an external stage, the maximum file size supported is 5 GB. This limit ensures efficient data transfer and management within Snowflake??s architecture

NEW QUESTION 154
- (Topic 3)
What can a Snowflake user do in the Admin area of Snowsight?

A. Analyze query performance.
B. Write queries and execute them.
C. Provide an overview of the listings in the Snowflake Marketplace.
D. Connect to Snowflake partners to explore extended functionality.

Answer: A

Explanation:
In the Admin area of Snowsight, users can analyze query performance, manage Snowflake warehouses, set up and view details about resource monitors, manage users and roles, and administer Snowflake accounts in their organization2.

NEW QUESTION 158
- (Topic 3)
Which of the following are handled by the cloud services layer of the Snowflake architecture? (Choose two.)

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

Answer: DE

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

NEW QUESTION 162
- (Topic 3)
Which of the following describes the Snowflake Cloud Services layer?

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

Answer: A

Explanation:
The Snowflake Cloud Services layer coordinates activities within the Snowflake account. It is responsible for tasks such as authentication, infrastructure management, metadata management, query parsing and optimization, and access control. References: Based on general cloud database architecture knowledge.

NEW QUESTION 164
- (Topic 3)
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: CD

Explanation:
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 168
- (Topic 3)
Which Snowflake tool would be BEST to troubleshoot network connectivity?

A. SnowCLI
B. SnowUI
C. SnowSQL
D. SnowCD

Answer: D

Explanation:
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%20Diagnosti c,their%20network%20connection%20to%20Snowflake.

NEW QUESTION 173
- (Topic 3)
Which REST API can be used with unstructured data?

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

Answer: C

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

NEW QUESTION 177
- (Topic 4)
Which Snowflake feature allows a user to track sensitive data for compliance, discovery, protection, and resource usage?

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

Answer: A

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

NEW QUESTION 182
- (Topic 4)
What tasks can an account administrator perform in the Data Exchange? (Select TWO).

A. Add and remove members.
B. Delete data categories.
C. Approve and deny listing approval requests.
D. Transfer listing ownership.
E. Transfer ownership of a provider profile.

Answer: AC

Explanation:
An account administrator in the Data Exchange can perform tasks such as adding and removing members and approving or denying listing approval requests. These tasks are part of managing the Data Exchange and ensuring that only authorized listings and members are part of it12.

NEW QUESTION 184
- (Topic 4)
What metadata does Snowflake store for rows in micro-partitions? (Select TWO).

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

Answer: AB

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

NEW QUESTION 186
- (Topic 4)
Which type of loop requires a BREAK statement to stop executing?

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

Answer: B

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

NEW QUESTION 189
- (Topic 4)
What factors impact storage costs in Snowflake? (Select TWO).

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

Answer: AC

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

NEW QUESTION 190

- (Topic 4)
What is the minimum Snowflake Edition that supports secure storage of Protected Health Information (PHI) data?

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

Answer: C

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

NEW QUESTION 193
- (Topic 4)
What is the primary purpose of a directory table in Snowflake?

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

Answer: D

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

NEW QUESTION 194
- (Topic 4)
Which parameter can be set at the account level to set the minimum number of days for which Snowflake retains historical data in Time Travel?

A. DATA_RETENTION_TIME_IN_DAYS
B. MAX_DATA_EXTENSION_TIME_IN_DAYS
C. MIN_DATA_RETENTION_TIME_IN_DAYS
D. MAX CONCURRENCY LEVEL

Answer: A

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

NEW QUESTION 198
- (Topic 4)
What is a directory table in Snowflake?

A. A separate database object that is used to store file-level metadata
B. An object layered on a stage that is used to store file-level metadata
C. A database object with grantable privileges for unstructured data tasks
D. A Snowflake table specifically designed for storing unstructured files

Answer: B

Explanation:
A directory table in Snowflake is an object layered on a stage that is used to store file-level metadata. It is not a separate database object but is conceptually similar to an external table because it stores metadata about the data files in the stage5.

NEW QUESTION 200
- (Topic 4)
What does the LATERAL modifier for the FLATTEN function do?

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

Answer: C

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

NEW QUESTION 203
- (Topic 4)
Who can grant object privileges in a regular schema?

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

Answer: A

Explanation:
In a regular schema within Snowflake, the object owner has the privilege to grant object privileges. The object owner is typically the role that created the object or to whom the ownership of the object has been transferred78.
References = [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 205
- (Topic 4)
Which command is used to start configuring Snowflake for Single Sign-On (SSO)?

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

Answer: C

Explanation:
To start configuring Snowflake for Single Sign-On (SSO), the CREATE SECURITY INTEGRATION command is used. This command sets up a security integration object in Snowflake, which is necessary for enabling SSO with external identity providers using SAML 2.01.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 206
- (Topic 4)
Which privilege must be granted by one role to another role, and cannot be revoked?

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

Answer: C

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

NEW QUESTION 210
- (Topic 4)
A column named "Data" contains VARIANT data and stores values as follows:

```
{
"Employee" :{
        id : 100,
        name : John,
        location : "New York"
        }
}
```

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: D

Explanation:
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 214
- (Topic 4)
When using the ALLOW CLIENT_MFA_CACHING parameter, how long is a cached Multi- Factor Authentication (MFA) token valid for?

A. 1 hour
B. 2 hours
C. 4 hours

D. 8 hours

Answer: C

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

NEW QUESTION 215
- (Topic 4)
Which Snowflake data types can be used to build nested hierarchical data? (Select TWO)

A. INTEGER
B. OBJECT
C. VARIANT
D. VARCHAR
E. LIST

Answer: BC

Explanation:
The Snowflake data types that can be used to build nested hierarchical data are OBJECT and VARIANT. These data types support the storage and querying of semi- structured data, allowing for the creation of complex, nested data structures

NEW QUESTION 217
- (Topic 4)
Which command is used to unload data from a Snowflake database table into one or more files in a Snowflake stage?

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

Answer: C

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

NEW QUESTION 221
- (Topic 4)
Which statements describe benefits of Snowflake's separation of compute and storage?
(Select TWO).

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

Answer: AE

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

NEW QUESTION 223
- (Topic 4)
Which Snowflake view is used to support compliance auditing?

A. ACCESS_HISTORY
B. COPY_HISTORY
C. QUERY_HISTORY
D. ROW ACCESS POLICIES

Answer: A

Explanation:
The ACCESS_HISTORY view in Snowflake is utilized to support compliance auditing. It provides detailed information on data access within Snowflake, including reads and writes by user queries. This view is essential for regulatory compliance auditing as it offers insights into the usage of tables and columns, and maintains a direct link between the user, the query, and the accessed data1.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 227
- (Topic 4)
Which statistics are displayed in a Query Profile that indicate that intermediate results do not fit in memory? (Select TWO).

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

Answer: CD

Explanation:
The Query Profile statistics that indicate intermediate results do not fit in memory are the bytes spilled to local storage and bytes spilled to remote storage2.

NEW QUESTION 232
- (Topic 4)
What is the purpose of the STRIP NULL_VALUES file format option when loading semi- structured data files into Snowflake?

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

Answer: D

Explanation:
The STRIP NULL_VALUES file format option, when set to TRUE, removes object or array elements that contain null values during the loading process of semi-structured data files into Snowflake. This ensures that the data loaded into Snowflake tables does not contain these null elements, which can be useful when the ??null?? values in files indicate missing values and have no other special meaning2.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 237
- (Topic 5)
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: A

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

NEW QUESTION 239
- (Topic 5)
Which Snowflow object does not consume and storage costs?

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

Answer: C

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

NEW QUESTION 241
- (Topic 5)
Which view can be used to determine if a table has frequent row updates or deletes?

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

Answer: B

Explanation:
The TABLE_STORAGE_METRICS view can be used to determine if a table has frequent row updates or deletes. This view provides detailed metrics on the storage utilization of tables within Snowflake, including metrics that reflect the impact of DML operations such as updates and deletes on table storage. For

example, metrics related to the number of active and deleted rows can help identify tables that experience high levels of row modifications, indicating frequent updates or deletions.
References:
? Snowflake Documentation: TABLE_STORAGE_METRICS View

NEW QUESTION 242
- (Topic 5)
Which use case does the search optimization service support?

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

Answer: D

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

NEW QUESTION 246
- (Topic 5)
Which function will provide the proxy information needed to protect Snowsight?

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

Answer: B

Explanation:
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 248
- (Topic 5)
How does a Snowflake stored procedure compare to a User-Defined Function (UDF)?

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

Answer: B

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

NEW QUESTION 253
- (Topic 5)
Which data types optimally store semi-structured data? (Select TWO).

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

Answer: AE

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

These data types provide the necessary flexibility to store and query semi-structured data efficiently in Snowflake.
References:
? Snowflake Documentation: Semi-structured Data Types

NEW QUESTION 257
- (Topic 5)
Which privilege is required to use the search optimization service in Snowflake?

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

Answer: C

Explanation:
To utilize the search optimization service in Snowflake, the correct syntax for granting privileges to a role involves specific commands that include adding search optimization capabilities:
? Option C: GRANT ADD SEARCH OPTIMIZATION ON SCHEMA <schema_name>
TO ROLE <role>. This command grants the specified role the ability to implement search optimization at the schema level, which is essential for enhancing search capabilities within that schema.
Options A and B do not include the correct verb "ADD," which is necessary for this specific type of grant command in Snowflake. Option D incorrectly mentions the database level, as search optimization privileges are typically configured at the schema level, not the database level.References:Snowflake documentation on the use of GRANT statements for configuring search optimization.

NEW QUESTION 260
- (Topic 5)
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: AC

Explanation:
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 ARRAYhttps://docs.snowflake.com/en/user-guide/tables-clustering-keys

NEW QUESTION 262
- (Topic 5)
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: D

Explanation:
To get the most detailed information about individual table storage details in Snowflake, theTABLE STORAGE METRICSview 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 263
- (Topic 5)
Which command can be used to list all the file formats for which a user has access privileges?

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

Answer: D

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

NEW QUESTION 264

- (Topic 5)
What is the MAXIMUM number of clusters that can be provisioned with a multi-cluster virtual warehouse?

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

Answer: C

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

NEW QUESTION 268
- (Topic 5)
User1, who has the SYSADMIN role, executed a query on Snowsight. User2, who is in the same Snowflake account, wants to view the result set of the query executed by User1 using the Snowsight query history.
What will happen if User2 tries to access the query history?

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

Answer: C

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

NEW QUESTION 270
- (Topic 5)
Which service or feature in Snowflake is used to improve the performance of certain types of lookup and analytical queries that use an extensive set of WHERE conditions?

A. Data classification
B. Query acceleration service
C. Search optimization service
D. Tagging

Answer: C

Explanation:
The Search Optimization Service in Snowflake is designed to improve the performance of specific types of queries, particularly those involving extensive sets of WHERE conditions. By maintaining a search index on tables, this service can accelerate lookup and analytical queries, making it a valuable feature for optimizing query performance and reducing execution times for complex searches.
References:
? Snowflake Documentation: Search Optimization Service

NEW QUESTION 272
- (Topic 5)
Which command removes a role from another role or a user in Snowflak?

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

Answer: B

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

NEW QUESTION 273
- (Topic 5)
What is the MINIMUM permission needed to access a file URL from an external stage?

A. MODIFY
B. READ
C. SELECT

D. USAGE

Answer: D

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

NEW QUESTION 276
- (Topic 6)
Which Snowflake objects can be restored using Time Travel? (Select VNO).

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

Answer: CD

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

NEW QUESTION 278
- (Topic 6)
What Snowflake recommendation is designed to ensure that staged data is only loaded once"?

A. Partitioning staged data files
B. Loading only the most recently-staged data files
C. Removing data files after loading
D. Identifying and removing duplicates after each data load

Answer: C

Explanation:
Snowflake recommends removing data files from the staging area after they have been loaded into the target table. This practice ensures that the data is only loaded once and prevents accidental reloading of the same data. By removing the files, you eliminate the risk of duplicate data loads.
? Stage the Data:Upload the data files to a Snowflake stage (internal or external).
? Load the Data:Use theCOPY INTOcommand to load the data from the stage into the Snowflake table.
? Remove the Data Files:After successfully loading the data, remove the data files from the stage using theREMOVEcommand.
References:
? Snowflake Documentation: Loading Data into Snowflake
? Snowflake Documentation: Staging Data Files
? Snowflake Documentation: COPY INTO Command

NEW QUESTION 283
- (Topic 6)
Which security models are used in Snowflake to manage access control? (Select TWO).

A. Discretionary Access Control (DAC)
B. Identity Access Management (1AM)
C. Mandatory Access Control (MAC)
D. Role-Based Access Control (RBAC)
E. Security Assertion Markup Language (SAML)

Answer: AD

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

NEW QUESTION 288
- (Topic 6)
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: A

Explanation:
The correct function to find the query ID of the second query executed in the current session isSELECT LAST_QUERY_ID(-2). TheLAST_QUERY_IDfunction 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-2would reference the second most recent query executed.
References:
? There's a clarification needed here; Snowflake's documentation indicatesLAST_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 289
- (Topic 6)
Which service or tool is a Command Line Interface (CLI) client used for connecting to Snowflake to execute SQL queries?

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

Answer: D

Explanation:
SnowSQL is the Command Line Interface (CLI) client provided by Snowflake for executing SQL queries and performing various tasks. It allows users to connect to their Snowflake accounts and interact with the Snowflake data warehouse.
? Installation:SnowSQL can be downloaded and installed on various operating systems.
? Configuration:Users need to configure SnowSQL with their Snowflake account credentials.
? Usage:Once configured, users can run SQL queries, manage data, and perform administrative tasks through the CLI.
References:
? Snowflake Documentation: SnowSQL
? Snowflake Documentation: Installing SnowSQL

NEW QUESTION 292
- (Topic 6)
Which data type can be used for floating-point numbers without losing precision?

A. BINARY
B. VARIANT
C. INTEGER
D. DOUBLE

Answer: D

Explanation:
In Snowflake, theDOUBLEdata type is used for floating-point numbers and can represent a wide range of values without losing precision. This data type is ideal for storing numerical values that require decimal precision.
? Define the Column:When creating a table, specify the column with theDOUBLEdata type to store floating-point numbers. CREATE TABLE example_table (
id INTEGER,
value DOUBLE
);
? Insert Data:Insert floating-point numbers into theDOUBLEcolumn. INSERT INTO example_table (id, value) VALUES (1, 123.456); References:
? Snowflake Documentation: Data Types
? Snowflake Documentation: Numeric Data Types

NEW QUESTION 296
- (Topic 6)
While unloading data into a stage, how can the user ensure that the output will be a single file?

A. Use the copy option files=single.
B. Use the COPY Option SINGLE=TRUE .
C. Use the get option SINGLE-TRUE.
D. Use the GET option FILES-SINGLE.

Answer: B

Explanation:
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.
References:
? Snowflake Documentation: COPY INTO <location>

NEW QUESTION 301
- (Topic 6)
A user needs to MINIMIZE the cost of large tables that are used to store transitory data. The data does not need to be protected against failures, because the data can be reconstructed outside of Snowflake.
What table type should be used?

A. Permanent
B. Transient
C. Temporary
D. Externa

Answer: B

Explanation:
For minimizing the cost of large tables that are used to store transitory data, which does not need to be protected against failures because it can be reconstructed outside of Snowflake, the best table type to use isTransient. Transient tables in Snowflake are designed for temporary or transitory data storage and offer reduced storage costs compared to permanent tables. However, unlike temporary tables, they persist across sessions until explicitly dropped.
? Why Transient Tables:Transient tables provide a cost-effective solution for storing data that is temporary but needs to be available longer than a single session. They have lower data storage costs because Snowflake does not maintain historical data (Time Travel) for as long as it does for permanent tables.
? Creating a Transient Table:
CREATETRANSIENTTABLEmy_transient_table (...);
? Use Case Considerations:Transient tables are ideal for scenarios where the data is not critical, can be easily recreated, and where cost optimization is a priority. They are suitable for development, testing, or staging environments where data longevity is not a concern.
Reference:For more details on transient tables and their usage scenarios, refer to the Snowflake documentation on table types: https://docs.snowflake.com/en/sql-reference/sql/create-table.html#table-types

NEW QUESTION 305
- (Topic 6)
How does the Access_History view enhance overall data governance pertaining to read and write operations? (Select TWO).

A. Shows how the accessed data was moved from the source lo the target objects
B. Provides a unified picture of what data was accessed and when it was accessed
C. Protects sensitive data from unauthorized access while allowing authorized users to access it at query runtime
D. Identifies columns with personal information and tags them so masking policies can be applied to protect sensitive data
E. Determines whether a given row in a table can be accessed by the user by filtering the data based on a given policy

Answer: BE

Explanation:
TheACCESS_HISTORYview in Snowflake is a powerful tool for enhancing data governance, especially concerning monitoring and auditing data access patterns for both read and write operations. The key ways in whichACCESS_HISTORYenhances overall data governance are:
? B. Provides a unified picture of what data was accessed and when it was accessed:This view logs details about query executions, including the objects (tables, views) accessed and the timestamps of these accesses. It's instrumental in auditing and compliance scenarios, where understanding the access patterns to sensitive data is critical.
? E. Determines whether a given row in a table can be accessed by the user by filtering the data based on a given policy:While this option is a bit of a misinterpretation of whatACCESS_HISTORYdirectly offers, it indirectly supports data governance by providing the information necessary to analyze access patterns. This analysis can then inform policy decisions, such as implementing Row-Level Security (RLS) to restrict access to specific rows based on user roles or attributes.
ACCESS_HISTORYdoes not automatically apply data masking or tag columns with personal information. However, the insights derived from analyzingACCESS_HISTORY can be used to identify sensitive data and inform the application of masking policies or other security measures.
References:
? Snowflake Documentation onACCESS_HISTORY: Access History

NEW QUESTION 310
- (Topic 6)
Which Snowflake table supports unstructured data?

A. Directory
B. Transient
C. Temporary
D. Permanent

Answer: D

Explanation:
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 313
- (Topic 6)
Which role is responsible for managing the billing and credit data within Snowflake?

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

Answer: A

Explanation:
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.
References:
? Snowflake Documentation: Account and Organization Roles

NEW QUESTION 314
- (Topic 6)

Which function unloads data from a relational table to JSON?

A. TRUNC<ID_NUMBER, -6)
B. TRUNC(ID_NUMBER, 5)
C. ID_NUMBER*100
D. TO_CHAR<ID NUMBER)

Answer: D

Explanation:
To unload data from a relational table to JSON format, you can use the TO_CHAR function. This function converts a number to a character string, which can then be serialized into JSON format. While there isn't a direct function specifically named for unloading to JSON, converting the necessary fields to a string representation is a common step in preparing data for JSON serialization.
References:
? Snowflake Documentation: TO_CHAR Function

NEW QUESTION 315
- (Topic 6)
Which access control entity in Snowflake can be created as part of a hierarchy within an account?

A. Securable object
B. Role
C. Privilege
D. User

Answer: B

Explanation:
In Snowflake, a role is an access control entity that can be created as part of a hierarchy within an account. Roles are used to grant and manage privileges in a structured and scalable manner.
? Understanding Roles:
? Role Hierarchy:
? Creating Roles:
? Example Usage: CREATE ROLE role1; CREATE ROLE role2;
GRANT ROLE role1 TO role2;
? uk.co.certification.simulator.questionpool.PList@546f4af7
? Benefits:
References:
? Snowflake Documentation: Access Control in Snowflake
? Snowflake Documentation: Creating and Managing Roles

NEW QUESTION 320
- (Topic 6)
When snaring data in Snowflake. what privileges does a Provider need to grant along with a share? (Select TWO).

A. USAGE on the specific tables in the database.
B. USAGE on the specific tables in the database.
C. MODIFY on 1Mb specific tables in the database.
D. USAGE on the database and the schema containing the tables to share
E. OPEBATE on the database and the schema containing the tables to share.

Answer: AD

Explanation:
When sharing data in Snowflake, the provider needs to grant the following privileges along with a share:
? A. USAGE on the specific tables in the database:This privilege allows the
consumers of the share to access the specific tables included in the share.
? D. USAGE on the database and the schema containing the tables to share:This privilege is necessary for the consumers to access the database and schema levels, enabling them to access the tables within those schemas.
These privileges are crucial for setting up secure and controlled access to the shared data, ensuring that only authorized users can access the specified resources.
Reference to Snowflake documentation on sharing data and managing access:
? Data Sharing Overview
? Privileges Required for Sharing Data

NEW QUESTION 322
- (Topic 6)
What takes the highest precedence in Snowflake file format options, when specified in multiple locations during data loading?

A. The stage definition
B. The table definition
C. The use of a copy into <table> statement
D. The use of a copy INTO <location> statement

Answer: C

Explanation:
When loading data into Snowflake, the file format options specified in theCOPY INTO <table>statement take the highest precedence over other locations such as the stage or table definitions. This ensures that any specific settings for a particular load operation are applied correctly.
? File Format Hierarchy:
? Example Usage: COPY INTO my_table FROM @my_stage

FILE_FORMAT = (FORMAT_NAME = 'my_format' FIELD_OPTIONALLY_ENCLOSED_BY
= '"');
References:
? Snowflake Documentation: Copy into Table
? Snowflake Documentation: File Format Options

NEW QUESTION 324
- (Topic 6)
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: C

Explanation:
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 329
- (Topic 6)
Which user preferences can be set for a user profile in Snowsight? (Select TWO).

A. Multi-Factor Authentication (MFA)
B. Default database
C. Default schema
D. Notification
E. Username

Answer: BC

Explanation:
In Snowsight, Snowflake's web interface, user preferences can be customized to enhance the user experience. Among these preferences, users can set a default database and default schema. These settings streamline the user experience by automatically selecting the specified database and schema when the user initiates a new session or query, reducing the need to manually specify these parameters for each operation. This feature is particularly useful for users who frequently work within a specific database or schema context.References:Snowflake Documentation on Snowsight User Preferences

NEW QUESTION 332
- (Topic 6)
Which function returns an integer between 0 and 100 when used to calculate the similarity of two strings?

A. APPROXIMATE_SIMILARITY
B. JAROWINKLER_SIMILARITY
C. APPROXIMATE_JACCARD_INDEX
D. MINHASH COMBINE

Answer: B

Explanation:
TheJAROWINKLER_SIMILARITYfunction in Snowflake returns an integer between 0 and 100, indicating the similarity of two strings based on the Jaro-Winkler similarity algorithm. This function is useful for comparing strings and determining how closely they match each other.
? Understanding JAROWINKLER_SIMILARITY:The Jaro-Winkler similarity metric is
a measure of similarity between two strings. The score is a number between 0 and 100, where 100 indicates an exact match and lower scores indicate less similarity.
? Usage Example:To compare two strings and get their similarity score, you can
use: SELECTJAROWINKLER_SIMILARITY('string1','string2')ASsimilarity_score;
? Application Scenarios:This function is particularly useful in data cleaning,
matching, and deduplication tasks where you need to identify similar but not identical strings, such as names, addresses, or product titles.
Reference:For more detailed information on theJAROWINKLER_SIMILARITYfunction and its usage, refer to the Snowflake documentation on string functions:
https://docs.snowflake.com/en/sql-reference/functions/jarowinkler_similarity.html

NEW QUESTION 337
- (Topic 6)
Which table function should be used to view details on a Directed Acyclic Graphic (DAG) run that is presently scheduled or is executing?

A. TASK_HISTORY
B. TASK_DEPENDENTS
C. CURRENT_TASK_GRAPHS
D. COMPLETE_TASK_GRAPHS

Answer: C

Explanation:
TheCURRENT_TASK_GRAPHStable function is designed to provide information on Directed Acyclic Graphs (DAGs) that are currently scheduled or executing within Snowflake. This function offers insights into the structure and status of task chains, enabling users to monitor and troubleshoot task executions. DAGs in

Snowflake represent sequences of tasks with dependencies, and understanding their current state is crucial for managing complex workflows.References:Snowflake Documentation on Task Management

NEW QUESTION 341
- (Topic 6)
The following settings are configured:
For how many days will the data be retained at the object level?

A. 2
B. 3
C. 5
D. 7

Answer: A

Explanation:
The settings shown in the image indicate that the data retention time in days is configured at two different levels: the account level and the object level. At the account level, theMIN_DATA_RETENTION_TIME_IN_DAYSis set to 5 days, and at the object level, theDATA_RETENTION_TIME_IN_DAYSis set to 2 days. Since the object level setting has a lower value, it takes precedence over the account level setting for the specific object. Therefore, the data will be retained for 2 days at the object level.References:Snowflake Documentation on Data Retention Policies

NEW QUESTION 343
- (Topic 6)
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 forinformation schema views, depending on the view.
E. Information schema views are read-only but account usage views are not.

Answer: CD

Explanation:
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 daysto a maximum of 6 months, depending on the specific view.References:Snowflake Documentation on Account Usage and Information Schema

NEW QUESTION 345
- (Topic 6)
What are potential impacts of storing non-native values like dates and timestamps in a variant column in Snowflake?

A. Faster query performance and increased storage consumption
B. Slower query performance and increased storage consumption
C. Faster query performance and decreased storage consumption
D. Slower query performance and decreased storage consumption

Answer: B

Explanation:
Storing non-native values, such as dates and timestamps, in a VARIANT column in Snowflake can lead to slower query performance and increased storage consumption. VARIANT is a semi-structured data type that allows storing JSON, AVRO, ORC, Parquet, or XML data in a single column. When non-native data types are stored as VARIANT, Snowflake must perform implicit conversion to process these values, which can slow down query execution. Additionally, because the VARIANT data type is designed to accommodate a wide variety of data formats, it often requires more storage spacecompared to storing data in native, strongly-typed columns that are optimized for specific data types.
The performance impact arises from the need to parse and interpret the semi-structured data on the fly during query execution, as opposed to directly accessing and operating on optimally stored data in its native format. Furthermore, the increased storage consumption is a result of the overhead associated with storing data in a format that is less space- efficient than the native formats optimized for specific types of data.
References:
? Snowflake Documentation on Semi-Structured Data: Semi-Structured Data

NEW QUESTION 346
- (Topic 6)
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. II 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: B

Explanation:
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 348
- (Topic 6)
Which Snowflake object can be used to record DML changes made to a table?

A. Snowpipe
B. Stage
C. Stream
D. Task

Answer: C

Explanation:
Snowflake Streams are used to track and record Data Manipulation Language (DML) changes made to a table. Streams capture changes such as inserts, updates, and deletes, which can then be processed by other Snowflake objects or external applications.
? Creating a Stream:
CREATE OR REPLACE STREAM my_stream ON TABLE my_table;
? Using Streams:Streams provide a way to process changes incrementally, making it easier to build efficient data pipelines.
? Consuming Stream Data:The captured changes can be consumed using SQL queries or Snowflake tasks.
References:
? Snowflake Documentation: Using Streams
? Snowflake Documentation: Change Data Capture (CDC) with Streams

NEW QUESTION 349
- (Topic 6)
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: BD

Explanation:
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 351
- (Topic 6)
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: C

Explanation:
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 353
- (Topic 6)
What is the benefit of using the STRIP_OUTER_ARRAY parameter with the COPY INTO
<table> command when loading data from a JSON file into a table?

A. It flattens multiple arrays into a single array.
B. It removes the outer array structure and loads separate rows of data
C. It transforms a pivoted table into an array.
D. It tokenizes each data string using the defined delimiters.

Answer: B

Explanation:
The STRIP_OUTER_ARRAY parameter in the COPY INTO <table> command is used when loading data from a JSON file into a table. This parameter removes the outer array structure from the JSON data and loads separate rows of data into the table.
? Understanding the STRIP_OUTER_ARRAY Parameter:

? How It Works:
? Example Usage: FROM @my_stage/file.json
FILE_FORMAT = (TYPE = 'JSON' STRIP_OUTER_ARRAY = TRUE);
? uk.co.certification.simulator.questionpool.PList@6ef05806
? Benefits:
References:
? Snowflake Documentation: COPY INTO <table>
? Snowflake Documentation: JSON File Format Options

NEW QUESTION 357
- (Topic 6)
How can a user MINIMIZE Continuous Data Protection costs when using large, high-churn, dimension tables?

A. Create transient tables and periodically copy them to permanent tables.
B. Create temporary tables and periodically copy them to permanent tables
C. Create regular tables with extended Time Travel and Fail-safe settings.
D. Create regular tables with default Time Travel and Fail-safe settings

Answer: A

Explanation:
To minimize Continuous Data Protection (CDP) costs when dealing with large, high-churn dimension tables in Snowflake, using transient tables is a recommended approach.
? Transient Tables: These are designed for data that does not require fail-safe
protection. They provide the benefit of reducing costs associated with continuous data protection since they do not have the seven-day Fail-safe period that is mandatory for permanent tables.
? Periodic Copying to Permanent Tables: By periodically copying data from transient
tables to permanent tables, you can achieve a balance between data protection and cost-efficiency. Permanent tables offer the extended data protection features, including Time Travel and Fail-safe, but these features can be applied selectively rather than continuously, reducing the overall CDP costs.
References:
? Snowflake Documentation on Transient Tables
? Snowflake Documentation on Time Travel & Fail-safe

NEW QUESTION 362
- (Topic 6)
Which type of role can be granted to a share?

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

Answer: B

Explanation:
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 365
- (Topic 6)
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: D

Explanation:
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 369

- (Topic 6)
What happens when a suspended virtual warehouse is resized in Snowflake?

A. It will return an error.
B. It will return a warning.
C. The suspended warehouse is resumed and new compute resources are provisioned immediately.
D. The additional compute resources are provisioned when the warehouse is resumed.

Answer: D

Explanation:
In Snowflake, resizing a virtual warehouse that is currently suspended does not immediately provision the new compute resources. Instead, the change in size is recorded, and the additional compute resources are provisioned when the warehouse is resumed. This means that the action of resizing a suspended warehouse does not cause it to resume operation automatically. The warehouse remains suspended until an explicit command to resume it is issued, or until it automatically resumes upon the next query execution that requires it.
This behavior allows for efficient management of compute resources, ensuring that credits are not consumed by a warehouse that is not in use, even if its size is adjusted while it is suspended.
Reference: Snowflake Documentation on Resizing Warehouses
(https://docs.snowflake.com/en/user-guide/warehouses-tasks.html#resizing-a-warehouse)

NEW QUESTION 372
- (Topic 6)
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: D

Explanation:
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 375
- (Topic 6)
Why would a Snowflake user load JSON data into a VARIANT column instead of a string column?

A. A VARIANT column is more secure than a string column
B. A VARIANT column compresses data and a string column does not.
C. A variant column can be used to create a data hierarchy and a string column cannot
D. A VARIANT column will have a better query performance than a string column.

Answer: C

Explanation:
A VARIANT column in Snowflake is specifically designed to store semi-structured data, such as JSON, and allows for the creation of a data hierarchy. Unlike string columns, VARIANT columns can natively handle JSON data structures, enabling complex querying and manipulation of hierarchical data using functions designed for semi-structured data. References:
? Snowflake Documentation: VARIANT Data

NEW QUESTION 377
- (Topic 6)
Which role must be used to create resource monitors?

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

Answer: B

Explanation:
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.
References:
? Snowflake Documentation: Resource Monitors

NEW QUESTION 380
- (Topic 6)
Which command can be used to list all network policies available in an account?

A. DESCRIBE SESSION POLICY

B. DESCRIBE NETWORK POLICY
C. SHOW SESSION POLICIES
D. SHOW NETWORK POLICIES

Answer: D

Explanation:
To list all network policies available in an account, the correct command is SHOW NETWORK POLICIES. Network policies in Snowflake are used to define and enforce rules for how users can connect to Snowflake, including IP whitelisting and other connection requirements. TheSHOW NETWORK POLICIEScommand provides a list of all network policies defined within the account, along with their details.
TheDESCRIBE SESSION POLICYandDESCRIBE NETWORK POLICYcommands do not
exist in Snowflake SQL syntax. TheSHOW SESSION POLICIEScommand is also incorrect, as it does not pertain to the correct naming convention used by Snowflake for network policy management.
UsingSHOW NETWORK POLICIESwithout any additional parameters will display all network policies in the account, which is useful for administrators to review and manage the security configurations pertaining to network access.
Reference: Snowflake Documentation on Network Policies
(https://docs.snowflake.com/en/sql-reference/sql/show-network-policies.html)

NEW QUESTION 385
- (Topic 6)
Which semi-structured file format is a compressed, efficient, columnar data representation?

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

Answer: D

Explanation:
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 386
- (Topic 6)
By default, which role can create resource monitors?

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

Answer: A

Explanation:
The role that can by default create resource monitors in Snowflake is theACCOUNTADMIN 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 theACCOUNTADMINrole. 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

NEW QUESTION 390
- (Topic 6)
What can the Snowflake SCIM API be used to manage? (Select TWO).

A. Integrations
B. Network policies
C. Session policies
D. Roles
E. Users

Answer: DE

Explanation:
The Snowflake SCIM (System for Cross-domain Identity Management) API is used for automated user and role management. It enables integration with identity providers (IdPs) for the provisioning and deprovisioning of user accounts and roles in Snowflake. This helps in managing access control and permissions systematically and aligns with identity governance practices.
References:
? Snowflake Documentation: Managing Users and Roles with SCIM API

NEW QUESTION 394
- (Topic 6)
How can a Snowflake user post-process the result of SHOW FILE FORMATS?

A. Use the RESULT_SCAN function.
B. Create a CURSOR for the command.
C. Put it in the FROM clause in brackets.
D. Assign the command to RESULTSET.

Answer: A

Explanation:
first run SHOW FILE FORMATS
then SELECT * FROM TABLE(RESULT_SCAN(LAST_QUERY_ID(-1)))
https://docs.snowflake.com/en/sql-reference/functions/result_scan#usage-notes

NEW QUESTION 396
- (Topic 6)
What compute resource is used when loading data using Snowpipe?

A. Snowpipe uses virtual warehouses provided by the user.
B. Snowpipe uses an Apache Kafka server for its compute resources.
C. Snowpipe uses compute resources provided by Snowflake.
D. Snowpipe uses cloud platform compute resources provided by the use

Answer: C

Explanation:
Snowpipe is Snowflake's continuous data ingestion service that allows for loading data as soon as it's available in a cloud storage stage. Snowpipe uses compute resources managed by Snowflake, separate from the virtual warehouses that users create for querying data. This means that Snowpipe operations do not consume the computational credits of user-created virtual warehouses, offering an efficient and cost-effective way to continuously load data into Snowflake.
References:
? Snowflake Documentation: Understanding Snowpipe

NEW QUESTION 398
- (Topic 6)
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: BE

Explanation:
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 401
- (Topic 6)
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: D

Explanation:
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 406
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

COF-C02 Practice Exam Features:

* COF-C02 Questions and Answers Updated Frequently

* COF-C02 Practice Questions Verified by Expert Senior Certified Staff

* COF-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* COF-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year