Vendor: Snowflake

Exam Code: COF-C02

Exam Name: SnowPro Core Certification Exam

Version: DEMO

QUESTION 1
What features does Snowflake Time Travel enable?

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

Answer: BC
Explanation:
Snowflake Time Travel is a powerful feature that allows users to access historical data within a defined period. It enables two key capabilities:
B . Restoring data-related objects that have been deleted within the past 90 days: Time Travel cn be used to restore tables, schemas, and databases that have been accidentally or intentionally deleted within the Time Travel retention period.
C. Conducting point-in-time analysis for BI reporting: It allows users to query historical data as it appeared at a specific point in time within the Time Travel retention period, which is crucial for business intelligence and reporting purposes.
While Time Travel does allow querying of past data, it is limited to the retention period set for the Snowflake account, which is typically 1 day for standard accounts and can be extended up to 90 days for enterprise accounts. It does not enable querying or restoring objects created or deleted beyond the retention period, nor does it provide analysis over all periods of time.
References:

QUESTION 2
Which statement about billing applies to Snowflake credits?

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

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

QUESTION 3
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.

QUESTION 4
When reviewing the load for a warehouse using the load monitoring chart, the chart indicates that a high volume of Queries are always queuing in the warehouse According to recommended best practice, what should be done to reduce the Queue volume? (Select TWO).

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

Answer: AB
Explanation:
To address a high volume of queries queuing in a warehouse, Snowflake recommends two best practices:
A. Use multi-clustered warehousing to scale out warehouse capacity: This approach allows for the distribution of queries across multiple clusters within a warehouse, effectively managing the load and reducing the queue volume.
B. Scale up the warehouse size to allow Queries to execute faster: Increasing the size of the warehouse provides more compute resources, which can reduce the time it takes for queries to execute and thus decrease the number of queries waiting in the queue. These strategies help to optimize the performance of the warehouse by ensuring that resources are scaled appropriately to meet demand.

QUESTION 5
Which of the following objects can be shared through secure data sharing?

A. Masking policy
B. Stored procedure
C. Task
D. External table

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

QUESTION 6
Which of the following commands cannot be used within a reader account?

A. CREATE SHARE
B. ALTER WAREHOUSE
C. DROP ROLE
D. SHOW SCHEMAS
E. DESCRBE TABLE

Answer: A
Explanation:
In Snowflake, a reader account is a type of account that is intended for consuming shared data rather than performing any data management or DDL operations. The CREATE SHARE command is used to share data from your account with another account, which is not a capability provided to reader accounts. Reader accounts are typically restricted from creating shares, as their primary purpose is to read shared data rather than to share it themselves.

QUESTION 7
A user unloaded a Snowflake table called mytable to an internal stage called mystage. Which command can be used to view the list of files that has been uploaded to the staged?

A. list @mytable;
B. list @%raytable;
C. list @ %m.ystage;
D. list @mystage;

Answer: D
Explanation:
The command list @mystage; is used to view the list of files that have been uploaded to an internal stage in Snowflake. The list command displays the metadata for all files in the specified stage, which in this case is mystage. This command is particularly useful for verifying that files have been successfully unloaded from a Snowflake table to the stage and for managing the files within the stage.

QUESTION 8
Which of the following Snowflake capabilities are available in all Snowflake editions? (Select TWO)

A. Customer-managed encryption keys through Tri-Secret Secure
B. Automatic encryption of all data
C. Up to 90 days of data recovery through Time Travel
D. Object-level access control
E. Column-level security to apply data masking policies to tables and views

Answer: BD
Explanation:
In all Snowflake editions, two key capabilities are universally available:
B. Automatic encryption of all data: Snowflake automatically encrypts all data stored in its platform, ensuring security and compliance with various regulations. This encryption is transparent to users and does not require any configuration or management.
D. Object-level access control: Snowflake provides granular access control mechanisms that allow administrators to define permissions at the object level, including databases, schemas,

tables, and views. This ensures that only authorized users can access specific data objects. These features are part of Snowflake's commitment to security and governance, and they are included in every edition of the Snowflake Data Cloud.

QUESTION 9
Which command is used to unload data from a Snowflake table into a file in a stage?

A. COPY INTO
B. GET
C. WRITE
D. EXTRACT INTO

Answer: A
Explanation:
The COPY INTO command is used in Snowflake to unload data from a table into a file in a stage. This command allows for the export of data from Snowflake tables into flat files, which can then be used for further analysis, processing, or storage in external systems.

QUESTION 10
How often are encryption keys automatically rotated by Snowflake?

A. 30 Days
B. 60 Days
C. 90 Days
D. 365 Days

Answer: A
Explanation:
Snowflake automatically rotates encryption keys when they are more than 30 days old. Active keys are retired, and new keys are created. This process is part of Snowflake's comprehensive security measures to ensure data protection and is managed entirely by the Snowflake service without requiring user intervention.

QUESTION 11
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.

QUESTION 12
A user has an application that writes a new Tile to a cloud storage location every 5 minutes. What would be the MOST efficient way to get the files into Snowflake?

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

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

QUESTION 13
Which of the following are best practice recommendations that should be considered when loading data into Snowflake? (Select TWO).

A. Load files that are approximately 25 MB or smaller.
B. Remove all dates and timestamps.
C. Load files that are approximately 100-250 MB (or larger)
D. Avoid using embedded characters such as commas for numeric data types
E. Remove semi-structured data types

Answer: CD
Explanation:
When loading data into Snowflake, it is recommended to:
C. Load files that are approximately 100-250 MB (or larger): This size is optimal for parallel processing and can help to maximize throughput. Smaller files can lead to overhead that outweighs the actual data processing time.
D. Avoid using embedded characters such as commas for numeric data types: Embedded characters can cause issues during data loading as they may be interpreted incorrectly. It's best to clean the data of such characters to ensure accurate and efficient data loading. These best practices are designed to optimize the data loading process, ensuring that data is loaded quickly and accurately into Snowflake.

QUESTION 14
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.

QUESTION 15
A user has unloaded data from Snowflake to a stage
Which SQL command should be used to validate which data was loaded into the stage?

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

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

QUESTION 16
What happens when a cloned table is replicated to a secondary database? (Select TWO)

A. A read-only copy of the cloned tables is stored.
B. The replication will not be successful.
C. The physical data is replicated
D. Additional costs for storage are charged to a secondary account
E. Metadata pointers to cloned tables are replicated

Answer: CE
Explanation:
When a cloned table is replicated to a secondary database in Snowflake, the following occurs:
C. The physical data is replicated: The actual data of the cloned table is physically replicated to the secondary database. This ensures that the secondary database has its own copy of the data, which can be used for read-only purposes or failover scenarios.
E. Metadata pointers to cloned tables are replicated: Along with the physical data, the metadata pointers that refer to the cloned tables are also replicated. This metadata includes information about the structure of the table and any associated properties. It's important to note that while the physical data and metadata are replicated, the secondary database is typically read-only and cannot be used for write operations. Additionally, while there may be additional storage costs associated with the secondary account, this is not a direct result of the replication process but

rather a consequence of storing additional data.

QUESTION 17
Which data types does Snowflake support when querying semi-structured data? (Select TWO)

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

Answer: AB
Explanation:
Snowflake supports querying semi-structured data using specific data types that are capable of handling the flexibility and structure of such data. The data types supported for this purpose are:
A. VARIANT: This is a universal data type that can store values of any other type, including structured and semi-structured types. It is particularly useful for handling JSON, Avro, ORC, Parquet, and XML data formats.
B. ARRAY: An array is a list of elements that can be of any data type, including VARIANT, and is used to handle semi-structured data that is naturally represented as a list. These data types are part of Snowflake's built-in support for semi-structured data, allowing for the storage, querying, and analysis of data that does not fit into the traditional row-column format.

QUESTION 18
Which of the following describes how multiple Snowflake accounts in a single organization relate to various cloud providers?

A. Each Snowflake account can be hosted in a different cloud vendor and region.
B. Each Snowflake account must be hosted in a different cloud vendor and region
C. All Snowflake accounts must be hosted in the same cloud vendor and region
D. Each Snowflake account can be hosted in a different cloud vendor, but must be in the same region.

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

QUESTION 19
A user is loading JSON documents composed of a huge array containing multiple records into Snowflake. The user enables the strip__outer_array file format option What does the STRIP_OUTER_ARRAY file format do?

A. It removes the last element of the outer array.
B. It removes the outer array structure and loads the records into separate table rows,
C. It removes the trailing spaces in the last element of the outer array and loads the records into

separate table columns

D. It removes the NULL elements from the JSON object eliminating invalid data and enables the ability to load the records

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

QUESTION 20
What are the default Time Travel and Fail-safe retention periods for transient tables?

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

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

QUESTION 21
What is a best practice after creating a custom role?

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

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

QUESTION 22
Which of the following Snowflake objects can be shared using a secure share? (Select TWO).

A. Materialized views
B. Sequences
C. Procedures
D. Tables

E. Secure User Defined Functions (UDFs)

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

QUESTION 23
Will data cached in a warehouse be lost when the warehouse is resized?

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

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

QUESTION 24
Which Snowflake partner specializes in data catalog solutions?

A. Alation
B. DataRobot
C. dbt
D. Tableau

Answer: A
Explanation:
Alation is known for specializing in data catalog solutions and is a partner of Snowflake. Data catalog solutions are essential for organizations to effectively manage their metadata and make it easily accessible and understandable for users, which aligns with the capabilities provided by Alation.

QUESTION 25
What is the MOST performant file format for loading data in Snowflake?

A. CSV (Unzipped)
B. Parquet
C. CSV (Gzipped)
D. ORC

Answer: B
Explanation:
Parquet is a columnar storage file format that is optimized for performance in Snowflake. It is designed to be efficient for both storage and query performance, particularly for complex queries

on large datasets. Parquet files support efficient compression and encoding schemes, which can lead to significant savings in storage and speed in query processing, making it the most performant file format for loading data into Snowflake.

QUESTION 26
Which copy INTO command outputs the data into one file?

A. SINGLE=TRUE
B. MAX_FILE_NUMBER=1
C. FILE_NUMBER=1
D. MULTIPLE=FAISE

Answer: B
Explanation:
The COPY INTO command in Snowflake can be configured to output data into a single file by setting the MAX_FILE_NUMBER option to 1. This option limits the number of files generated by the command, ensuring that only one file is created regardless of the amount of data being exported.

QUESTION 27
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.

QUESTION 28
Which Snowflake technique can be used to improve the performance of a query?

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

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

Thank You for Trying Our Product

Braindump2go Certification Exam Features:

★ More than **99,900** Satisfied Customers Worldwide.

★ Average **99.9%** Success Rate.

★ **Free Update** to match latest and real exam scenarios.

★ **Instant Download** Access! No Setup required.

★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.

★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.

★ **100%** Guaranteed Success or **100%** Money Back Guarantee.

★ **Fast**, helpful support **24x7**.

View list of all certification exams: http://www.braindump2go.com/all-products.html

10% Discount Coupon Code: **ASTR14**