Snowflake

Exam Questions COF-C02

SnowPro Core Certification Exam (COF-C02)

NEW QUESTION 1
- (Topic 1)
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.
References:
? Snowflake Documentation on Security Features
? SnowPro® Core Certification Exam 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 a limitation of a Materialized View?

A. A Materialized View cannot support any aggregate functions
B. A Materialized View can only reference up to two tables
C. A Materialized View cannot be joined with other tables
D. A Materialized View cannot be defined with a JOIN

Answer: D

Explanation:
Materialized Views in Snowflake are designed to store the result of a query and can be refreshed to maintain up-to-date data. However, they have certain limitations, one of which is that they cannot be defined using a JOIN clause. This means that a Materialized View can only be created based on a single source table and cannot combine data from multiple tables using JOIN operations.
References:
? Snowflake Documentation on Materialized Views
? SnowPro® Core Certification Study Guide

NEW QUESTION 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)
Which of the following Snowflake features provide continuous data protection automatically? (Select TWO).

A. Internal stages
B. Incremental backups
C. Time Travel
D. Zero-copy clones
E. Fail-safe

Answer: CE

Explanation:
Snowflake??s Continuous Data Protection (CDP) encompasses a set of features that help protect data stored in Snowflake against human error, malicious acts, and software failure. Time Travel allows users to access historical data (i.e., data that has been changed or deleted) for a defined period, enabling querying and restoring of data. Fail-safe is an additional layer of data protection that provides a recovery option in the event of significant data loss or corruption, which can only be performed by Snowflake. References:
? Continuous Data Protection | Snowflake Documentation1
? Data Storage Considerations | Snowflake Documentation2
? Snowflake SnowPro Core Certification Study Guide3
? Snowflake Data Cloud Glossary
https://docs.snowflake.com/en/user-guide/data-availability.html

NEW QUESTION 6
- (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 7
- (Topic 1)
Which of the following are valid methods for authenticating users for access into Snowflake? (Select THREE)

A. SCIM
B. Federated authentication
C. TLS 1.2
D. Key-pair authentication
E. OAuth
F. OCSP authentication

Answer: BDE

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

NEW QUESTION 8
- (Topic 1)
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.
References:
? Snowflake Documentation on Snowpipe
? SnowPro® Core Certification Study Guide

NEW QUESTION 9
- (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 10
- (Topic 1)
In which scenarios would a user have to pay Cloud Services costs? (Select TWO).

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

Answer: AE

Explanation:
In Snowflake, Cloud Services costs are incurred when the Cloud Services usage exceeds 10% of the compute usage (measured in credits). Therefore, scenarios A and E would result in Cloud Services charges because the Cloud Services usage is more than 10% of the compute credits used.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake??s official documentation on billing and usage1

NEW QUESTION 10
- (Topic 1)
True or False: When you create a custom role, it is a best practice to immediately grant that role to ACCOUNTADMIN.

A. True
B. False

Answer: B

Explanation:
The ACCOUNTADMIN role is the most powerful role in Snowflake and should be limited to a select number of users within an organization. It is responsible for account-level configurations and should not be used for day-to-day object creation or management. Granting a custom role to ACCOUNTADMIN could inadvertently give broad access to users with this role, which is not a recommended security practice.
Reference:https://docs.snowflake.com/en/user-guide/security-access-control-considerations.html

NEW QUESTION 11
- (Topic 1)
What tasks can be completed using the copy command? (Select TWO)

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

Answer: CD

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

NEW QUESTION 16
- (Topic 1)
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.
References:
? Understanding Encryption Key Management in Snowflake

NEW QUESTION 17
- (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 20
- (Topic 1)
What are two ways to create and manage Data Shares in Snowflake? (Choose two.)

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

Answer: AC

Explanation:
In Snowflake, Data Shares can be created and managed in two primary ways:
? Via the Snowflake Web Interface (UI): Users can create and manage shares through the graphical interface provided by Snowflake, which allows for a user-friendly experience.
? Via SQL commands: Snowflake also allows the creation and management of
shares using SQL commands. This method is more suited for users who prefer scripting or need to automate the process.
Reference:https://docs.snowflake.com/en/user-guide/data-sharing-provider.html

NEW QUESTION 21
- (Topic 1)
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.
References:
? Snowflake Documentation on Stages
? SnowPro® Core Certification Study Guide

NEW QUESTION 23
- (Topic 1)
Which command can be used to load data into an internal stage?

A. LOAD
B. copy
C. GET
D. PUT

Answer: D

Explanation:
The PUT command is used to load data into an internal stage in Snowflake. This command uploads data files from a local file system to a named internal stage, making the data available for subsequent loading into a Snowflake table using the COPY INTO command.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide

? Snowflake Documentation on Data Loading

NEW QUESTION 26
- (Topic 1)
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. Ye
C. because the compute resource is replaced in its entirety with a new compute resource.
D. N
E. because the size of the cache is independent from the warehouse size
F. Ye
G. 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. References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Virtual Warehouse Performance1

NEW QUESTION 31
- (Topic 1)
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 records1.
References:
? Snowflake Documentation on JSON File Format
? [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 35
- (Topic 1)
A sales table FCT_SALES has 100 million records. The following Query was executed
SELECT COUNT (1) FROM FCT SALES;
How did Snowflake fulfill this query?

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

Answer: D

Explanation:
Snowflake is designed to optimize query performance by utilizing metadata for certain types of queries. When executing a COUNT query, Snowflake can often fulfill the request by accessing metadata about the table??s row count, rather than scanning the entire table or micro-partitions. This is particularly efficient for large tables like FCT_SALES with a significant number of records. The metadata layer maintains statistics about the table, including the row count, which enables Snowflake to quickly return the result of a COUNT query without the need to perform a full scan. References:
? Snowflake Documentation on Metadata Management
? SnowPro® Core Certification Study Guide

NEW QUESTION 37
- (Topic 1)
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.

References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Data Loading1

NEW QUESTION 38
- (Topic 1)
True or False: A 4X-Large Warehouse may, at times, take longer to provision than a X- Small Warehouse.

A. True
B. False

Answer: A

Explanation:
Provisioning time can vary based on the size of the warehouse. A 4X-Large Warehouse typically has more resources and may take longer to provision compared to a X-Small Warehouse, which has fewer resources and can generally be provisioned more quickly.References: Understanding and viewing Fail-safe | Snowflake Documentation

NEW QUESTION 43
- (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 47
- (Topic 1)
True or False: Fail-safe can be disabled within a Snowflake account.

A. True
B. False

Answer: B

Explanation:
Reference:https://docs.snowflake.com/en/user-guide/data-failsafe.html
Separate and distinct from Time Travel, Fail-safe ensures historical data is protected in the event of a system failure or other catastrophic event, e.g. a hardware failure or security breach. Fail-safe feature cannot be enabled or disabled from the user end.

NEW QUESTION 48
- (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 51
- (Topic 1)
In which use cases does Snowflake apply egress charges?

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

Answer: C

Explanation:
Snowflake applies egress charges in the case of database replication when data is transferred out of a Snowflake region to another region or cloud provider. This is because the data transfer incurs costs associated with moving data across different networks. Egress charges are not applied for data sharing within the same region, query result retrieval, or loading data into Snowflake, as these actions do not involve data transfer across regions.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Data Replication and Egress Charges1

NEW QUESTION 54
- (Topic 1)
Which of the following indicates that it may be appropriate to use a clustering key for a table? (Select TWO).

A. The table contains a column that has very low cardinality
B. DML statements that are being issued against the table are blocked
C. The table has a small number of micro-partitions
D. Queries on the table are running slower than expected
E. The clustering depth for the table is large

Answer: DE

Explanation:
A clustering key in Snowflake is used to co-locate similar data within the same micro-partitions to improve query performance, especially for large tables where data is not naturally ordered or has become fragmented due to extensive DML operations. The appropriate use of a clustering key canlead to improved scan efficiency and better column compression, resulting in faster query execution times.
The indicators that it may be appropriate to use a clustering key for a table include:
? D. Queries on the table are running slower than expected: This can happen when the data in the table is not well-clustered, leading to inefficient scans during query execution.
? E. The clustering depth for the table is large: A large clustering depth indicates that the table??s data is spread across many micro-partitions, which can degrade query performance as more data needs to be scanned.
References:
? Snowflake Documentation on Clustering Keys & Clustered Tables
? Snowflake Documentation on SYSTEM$CLUSTERING_INFORMATION
? Stack Overflow discussion on cluster key selection in Snowflake

NEW QUESTION 58
- (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 61
- (Topic 1)
When reviewing a query profile, what is a symptom that a query is too large to fit into the memory?

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

Answer: D

Explanation:
When a query in Snowflake is too large to fit into the available memory, it will start spilling to remote storage. This is an indication that the memory allocated for the query is insufficient for its execution, and as a result, Snowflake uses remote disk storage to handle the overflow. This spill to remote storage can lead to slower query performance due to the additional I/O operations required.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Query Profile1
? Snowpro Core Certification Exam Flashcards2

NEW QUESTION 62
- (Topic 1)
What happens when a virtual warehouse is resized?

A. When increasing the size of an active warehouse the compute resource for all running and queued queries on the warehouse are affected
B. When reducing the size of a warehouse the compute resources are removed only when they are no longer being used to execute any current statements.

C. The warehouse will be suspended while the new compute resource is provisioned and will resume automatically once provisioning is complete.
D. Users who are trying to use the warehouse will receive an error message until the resizing is complete

Answer: A

Explanation:
When a virtual warehouse in Snowflake is resized, specifically when it is increased in size, the additional compute resources become immediately available to all running and queued queries. Thismeans that the performance of these queries can improve due to the increased resources. Conversely, when the size of a warehouse is reduced, the compute resources are not removed until they are no longer being used by any current operations1.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Virtual Warehouses2

NEW QUESTION 63
- (Topic 1)
A marketing co-worker has requested the ability to change a warehouse size on their medium virtual warehouse called mktg WH.
Which of the following statements will accommodate this request?

A. ALLOW RESIZE ON WAREHOUSE MKTG WH TO USER MKTG LEAD;
B. GRANT MODIFY ON WAREHOUSE MKTG WH TO ROLE MARKETING;
C. GRANT MODIFY ON WAREHOUSE MKTG WH TO USER MKTG LEAD;
D. GRANT OPERATE ON WAREHOUSE MKTG WH TO ROLE MARKETING;

Answer: B

Explanation:
The correct statement to accommodate the request for a marketing co- worker to change the size of their medium virtual warehouse called mktg WH is to grant the MODIFY privilege on the warehouse to the ROLE MARKETING. This privilege allows the role to change the warehouse size among other properties.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Access Control Privileges1

NEW QUESTION 65
- (Topic 1)
During periods of warehouse contention which parameter controls the maximum length of time a warehouse will hold a query for processing?

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

Answer: B

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

NEW QUESTION 70
- (Topic 1)
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.
References:
? Snowflake Documentation on Stages
? SnowPro® Core Certification Study Guide

NEW QUESTION 73
- (Topic 1)
Which statement about billing applies to Snowflake credits?

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

Answer: D

Explanation:
Snowflake credits are the unit of measure for the compute resources used in Snowflake. The number of credits consumed depends on the size of the virtual warehouse and the time it is running. Larger warehouses consume more credits per hour than smaller ones, and credits are billed for the time the warehouse is active, regardless of the actual usage within that time.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 77
- (Topic 1)
What happens to the underlying table data when a CLUSTER BY clause is added to a Snowflake table?

A. Data is hashed by the cluster key to facilitate fast searches for common data values
B. Larger micro-partitions are created for common data values to reduce the number of partitions that must be scanned
C. Smaller micro-partitions are created for common data values to allow for more parallelism
D. Data may be colocated by the cluster key within the micro-partitions to improve pruning performance

Answer: D

Explanation:
When a CLUSTER BY clause is added to a Snowflake table, it specifies one or more columns to organize the data within the table??s micro-partitions. This clustering aims to colocate data with similar values in the same or adjacent micro-partitions. By doing so, it enhances the efficiency of query pruning, where the Snowflake query optimizer can skip over irrelevant micro-partitions that do not contain the data relevant to the query, thereby improving performance.
References:
? Snowflake Documentation on Clustering Keys & Clustered Tables1.
? Community discussions on how source data??s ordering affects a table with a cluster key

NEW QUESTION 78
- (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 80
- (Topic 1)
Which of the following describes how clustering keys work in Snowflake?

A. Clustering keys update the micro-partitions in place with a full sort, and impact the DML operations.
B. Clustering keys sort the designated columns over time, without blocking DML operations
C. Clustering keys create a distributed, parallel data structure of pointers to a table's rows and columns
D. Clustering keys establish a hashed key on each node of a virtual warehouse to optimize joins at run-time

Answer: B

Explanation:
Clustering keys in Snowflake work by sorting the designated columns over time. This process is done in the background and does not block data manipulation language (DML) operations, allowing for normal database operations to continue without interruption. The purpose of clustering keys is to organize the data within micro-partitions to optimize query performance1.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Clustering1

NEW QUESTION 82
- (Topic 1)
What is a responsibility of Snowflake's virtual warehouses?

A. Infrastructure management
B. Metadata management
C. Query execution
D. Query parsing and optimization
E. Management of the storage layer

Answer: C

Explanation:
The primary responsibility of Snowflake??s virtual warehouses is to execute queries. Virtual warehouses are one of the key components of Snowflake??s architecture, providing the compute power required to perform data processing tasks such as running SQL queries, performing joins, aggregations, and other data manipulations. References:

? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Virtual Warehouses1

NEW QUESTION 83
- (Topic 1)
When unloading to a stage, which of the following is a recommended practice or approach?

A. Set SINGLE: = true for larger files
B. Use OBJECT_CONSTRUCT (*) when using Parquet
C. Avoid the use of the CAST function
D. Define an individual file format

Answer: D

Explanation:
When unloading data to a stage, it is recommended to define an individual file format. This ensures that the data is unloaded in a consistent and expected format, which can be crucial for downstream processing and analysis2

NEW QUESTION 88
- (Topic 1)
A company's security audit requires generating a report listing all Snowflake logins (e.g.. date and user) within the last 90 days. Which of the following statements will return the required information?

A. SELECT LAST_SUCCESS_LOGIN, LOGIN_NAME FROM ACCOUNT_USAGE.USERS;
B. SELECT EVENT_TIMESTAMP, USER_NAMEFROM table(information_schema.login_history_by_user())
C. SELECT EVENT_TIMESTAMP, USER_NAME FROM ACCOUNT_USAGE.ACCESS_HISTORY;
D. SELECT EVENT_TIMESTAMP, USER_NAME FROM ACCOUNT_USAGE.LOGIN_HISTORY;

Answer: D

Explanation:
To generate a report listing all Snowflake logins within the last 90 days, the ACCOUNT_USAGE.LOGIN_HISTORY view should be used. This view provides information about login attempts, including successful and unsuccessful logins, and is suitable for security audits4.

NEW QUESTION 93
- (Topic 1)
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.
References:
? Snowflake Documentation on Data Loading Considerations
? [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 95
- (Topic 1)
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.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake??s official documentation and partner listings

NEW QUESTION 97
- (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 101
- (Topic 1)
What data is stored in the Snowflake storage layer? (Select TWO).

A. Snowflake parameters
B. Micro-partitions
C. Query history
D. Persisted query results
E. Standard and secure view results

Answer: BD

Explanation:
The Snowflake storage layer is responsible for storing data in an optimized, compressed, columnar format. This includes micro-partitions, which are the fundamental storage units that contain the actual data stored in Snowflake. Additionally, persisted query results, which are the results of queries that have been materialized and stored for future use, are also kept within this layer. This design allows for efficient data retrieval and management within the Snowflake architecture1.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Key Concepts & Architecture | Snowflake Documentation2

NEW QUESTION 103
- (Topic 1)
What Snowflake role must be granted for a user to create and manage accounts?

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

Answer: A

Explanation:
The ACCOUNTADMIN role is required for a user to create and manage accounts in Snowflake. This role has the highest level of privileges and is responsible for managing all aspects of the Snowflake account, including the ability to create and manage other user accounts1.
https://docs.snowflake.com/en/user-guide/security-access-control-considerations.html

NEW QUESTION 107
- (Topic 2)
How does Snowflake Fail-safe protect data in a permanent table?

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

Answer: C

Explanation:
Snowflake??s Fail-safe provides a 7-day period during which data in a permanent table may be recoverable, but only by Snowflake Support, not by user operations3.

NEW QUESTION 110
- (Topic 2)
What are the correct parameters for time travel and fail-safe in the Snowflake Enterprise Edition?

A. Default Time Travel Retention is set to 0 day
B. Maximum Time Travel Retention is 30 day
C. Fail Safe retention time is 1 day.
D. Default Time Travel Retention is set to 1 da

E. Maximum Time Travel Retention is 365 day
F. Fail Safe retention time is 7 days.
G. Default Time Travel Retention is set to 0 day
H. Maximum Time Travel Retention is 90 day
I. Fail Safe retention time is 7 days.
J. Default Time Travel Retention is set to 1 da
K. Maximum Time Travel Retention is 90 day
L. Fail Safe retention time is 7 days.
M. Default Time Travel Retention is set to 7 day
N. Maximum Time Travel Retention is 1 da
O. Fail Safe retention time is 90 days.
P. Default Time Travel Retention is set to 90 day
Q. Maximum Time Travel Retention is 7 day
R. Fail Safe retention time is 356 days.

Answer: D

Explanation:
In the Snowflake Enterprise Edition, the default Time Travel retention is set to 1 day, the maximum Time Travel retention can be set up to 90 days, and the Fail-safe retention time is 7 days3.

NEW QUESTION 115
- (Topic 2)
The is the minimum Fail-safe retention time period for transient tables?

A. 1 day
B. 7 days
C. 12 hours
D. 0 days

Answer: D

Explanation:
Transient tables in Snowflake have a minimum Fail-safe retention time period of 0 days. This means that once the Time Travel retention period ends, there is no additional Fail-safe period for transient tables

NEW QUESTION 119
- (Topic 2)
When publishing a Snowflake Data Marketplace listing into a remote region what should be taken into consideration? (Choose two.)

A. There is no need to have a Snowflake account in the target region, a share will be created for each user.
B. The listing is replicated into all selected regions automatically, the data is not.
C. The user must have the ORGADMIN role available in at least one account to link accounts for replication.
D. Shares attached to listings in remote regions can be viewed from any account in an organization.
E. For a standard listing the user can wait until the first customer requests the data before replicating it to the target region.

Answer: BC

Explanation:
When publishing a Snowflake Data Marketplace listing into a remote region, it??s important to note that while the listing is replicated into all selected regions automatically, the data itself is not.Therefore, the data must be replicated separately. Additionally, the user must have the ORGADMIN role in at least one account to manage the replication of accounts1.

NEW QUESTION 120
- (Topic 2)
Which file formats are supported for unloading data from Snowflake? (Choose two.)

A. Avro
B. JSON
C. ORC
D. XML
E. Delimited (CSV, TSV, etc.)

Answer: BE

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

NEW QUESTION 121
- (Topic 2)
Users are responsible for data storage costs until what occurs?

A. Data expires from Time Travel
B. Data expires from Fail-safe
C. Data is deleted from a table
D. Data is truncated from a table

Answer: B

Explanation:
Users are responsible for data storage costs in Snowflake until the data expires from the Fail-safe period. Fail-safe is the final stage in the data lifecycle, following Time Travel, and provides additional protection against accidental data loss. Once data exits the Fail-safe state, users are no longer billed for its storage

NEW QUESTION 122
- (Topic 2)
Which of the following statements apply to Snowflake in terms of security? (Choose two.)

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

Answer: AC

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

NEW QUESTION 127
- (Topic 2)
What versions of Snowflake should be used to manage compliance with Personal Identifiable Information (PII) requirements? (Choose two.)

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

Answer: BC

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

NEW QUESTION 128
- (Topic 2)
When loading data into Snowflake via Snowpipe what is the compressed file size recommendation?

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

Answer: B

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

NEW QUESTION 129
- (Topic 2)
The Snowflake cloud services layer is responsible for which tasks? (Choose two.)

A. Local disk caching
B. Authentication and access control
C. Metadata management
D. Query processing
E. Database storage

Answer: BC

Explanation:
The Snowflake cloud services layer is responsible for tasks such as authentication and access control, ensuring secure access to the platform, and metadata management, which involves organizing and maintaining information about the data stored in Snowflake56.

NEW QUESTION 132
- (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 136
- (Topic 2)
Which of the following objects are contained within a schema? (Choose two.)

A. Role
B. Stream
C. Warehouse
D. External table
E. User
F. Share

Answer: BD

Explanation:
In Snowflake, a schema is a logical grouping of database objects, which can include streams and external tables. A stream is an object that allows users to query data that has changed in specified tables or views, and an external table is a table that references data stored outside of Snowflake. Roles, warehouses, users, and shares are not contained within a schema. References: SHOW OBJECTS, Database, Schema, & Share DDL

NEW QUESTION 138
- (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 141
- (Topic 2)
What is the maximum Time Travel retention period for a temporary Snowflake table?

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

Answer: B

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

NEW QUESTION 144
- (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 148
- (Topic 2)
Which Snowflake feature allows a user to substitute a randomly generated identifier for sensitive data, in order to prevent unauthorized users access to the data, before loading it into Snowflake?

A. External Tokenization
B. External Tables

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

Answer: A

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

NEW QUESTION 151
- (Topic 2)
Which of the following describes a Snowflake stored procedure?

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

Answer: D

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

NEW QUESTION 153
- (Topic 2)
The Snowflake Cloud Data Platform is described as having which of the following architectures?

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

Answer: C

Explanation:
Snowflake??s architecture is described as a multi-cluster, shared data architecture. This design combines the simplicity of a shared-disk architecture with the performance and scale-out benefits of a shared-nothing architecture, using a central repository accessible from all compute nodes2.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

NEW QUESTION 156
- (Topic 2)
Which of the following are best practices for loading data into Snowflake? (Choose three.)

A. Aim to produce data files that are between 100 MB and 250 MB in size, compressed.
B. Load data from files in a cloud storage service in a different region or cloud platform from the service or region containing the Snowflake account, to save on cost.
C. Enclose fields that contain delimiter characters in single or double quotes.
D. Split large files into a greater number of smaller files to distribute the load among the compute resources in an active warehouse.
E. When planning which warehouse(s) to use for data loading, start with the largest warehouse possible.
F. Partition the staged data into large folders with random paths, allowing Snowflake to determine the best way to load each file.

Answer: ACD

Explanation:
Best practices for loading data into Snowflake include aiming for data file sizes between 100 MB and 250 MB when compressed, as this size is optimal for parallel processing and minimizes overhead. Enclosing fields with delimiter characters in quotes ensures proper field recognition during the load process. Splitting large files into smaller ones allows for better distribution of the load across compute resources, enhancing performance and efficiency.

NEW QUESTION 158
- (Topic 2)
A table needs to be loaded. The input data is in JSON format and is a concatenation of multiple JSON documents. The file size is 3 GB. A warehouse size small is being used. The following COPY INTO command was executed:
COPY INTO SAMPLE FROM @~/SAMPLE.JSON (TYPE=JSON)
The load failed with this error:
Max LOB size (16777216) exceeded, actual size of parsed column is 17894470. How can this issue be resolved?

A. Compress the file and load the compressed file.
B. Split the file into multiple files in the recommended size range (100 MB - 250 MB).
C. Use a larger-sized warehouse.
D. Set STRIP_OUTER_ARRAY=TRUE in the COPY INTO command.

Answer: B

Explanation:
The error ??Max LOB size (16777216) exceeded?? indicates that the size of the parsed column exceeds the maximum size allowed for a single column value in Snowflake, which is 16 MB. To resolve this issue, the file should be split into multiple smaller files that are within the recommended size range of 100 MB to 250 MB. This will ensure that each JSON document within the files is smaller than the maximum LOB size allowed. Compressing the file, using a larger-sized

warehouse, or setting STRIP_OUTER_ARRAY=TRUE will not resolve the issue of the column size exceeding the maximum allowed. References: COPY INTO Error during Structured Data Load: ??Max LOB size (16777216) exceeded????

NEW QUESTION 163
- (Topic 2)
What is the purpose of multi-cluster virtual warehouses?

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

Answer: C

Explanation:
Multi-cluster virtual warehouses in Snowflake are designed to manage user and query concurrency needs. They allow for the allocation of additional clusters of compute resources, either statically or dynamically, to handle increased loads and reduce or eliminate the queuing of concurrent queries2.
https://docs.snowflake.com/en/user-guide/warehouses-multicluster.html#:~:text=Multi%2Dcluster%20warehouses%20enable%20you,during%20peak%20and%20off%20hours.

NEW QUESTION 166
- (Topic 2)
What types of data listings are available in the Snowflake Data Marketplace? (Choose two.)

A. Reader
B. Consumer
C. Vendor
D. Standard
E. Personalized

Answer: CE

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

NEW QUESTION 168
- (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 171
- (Topic 2)
Which methods can be used to delete staged files from a Snowflake stage? (Choose two.)

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

Answer: CD

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

NEW QUESTION 174
- (Topic 2)
What is the maximum total Continuous Data Protection (CDP) charges incurred for a temporary table?

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

Answer: D

Explanation:
For a temporary table, the maximum total Continuous Data Protection (CDP) charges incurred are for the duration of the session in which the table was created, which does not exceed 24 hours2.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation2

NEW QUESTION 178
- (Topic 2)
When loading data into Snowflake, how should the data be organized?

A. Into single files with 100-250 MB of compressed data per file
B. Into single files with 1-100 MB of compressed data per file
C. Into files of maximum size of 1 GB of compressed data per file
D. Into files of maximum size of 4 GB of compressed data per file

Answer: A

Explanation:
When loading data into Snowflake, it is recommended to organize the data into single files with 100-250 MB of compressed data per file. This size range is optimal for parallel processing and can help in achieving better performance during data loading operations. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 182
- (Topic 2)
By default, which Snowflake role is required to create a share?

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

Answer: D

Explanation:
By default, the Snowflake role required to create a share is ACCOUNTADMIN (D). This role has the necessary privileges to perform administrative tasks, including creating shares for data sharing purposes

NEW QUESTION 183
- (Topic 2)
Which of the following is an example of an operation that can be completed without requiring compute, assuming no queries have been executed previously?

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

Answer: B

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

NEW QUESTION 188
- (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 189
- (Topic 2)
What are common issues found by using the Query Profile? (Choose two.)

A. Identifying queries that will likely run very slowly before executing them
B. Locating queries that consume a high amount of credits
C. Identifying logical issues with the queries
D. Identifying inefficient micro-partition pruning

E. Data spilling to a local or remote disk

Answer: DE

Explanation:
The Query Profile in Snowflake is used to identify performance issues with queries. Common issues that can be found using the Query Profile include identifying inefficient micro-partition pruning (D) and data spilling to a local or remote disk (E). Micro- partition pruning is related to the efficiency of query execution, and data spilling occurs when the memory is insufficient, causing the query to write data to disk, which can slow down the query performance1.

NEW QUESTION 193
- (Topic 2)
Which of the following significantly improves the performance of selective point lookup queries on a table?

A. Clustering
B. Materialized Views
C. Zero-copy Cloning
D. Search Optimization Service

Answer: D

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

NEW QUESTION 196
- (Topic 2)
What is the MINIMUM edition of Snowflake that is required to use a SCIM security integration?

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

Answer: D

Explanation:
The minimum edition of Snowflake required to use a SCIM security integration is the Enterprise Edition. SCIM integrations are used for automated management of user identities and groups, and this feature is available starting from the Enterprise Edition of Snowflake. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 199
- (Topic 2)
A user created a new worksheet within the Snowsight Ul and wants to share this with teammates
How can this worksheet be shared?

A. Create a zero-copy clone of the worksheet and grant permissions to teammates
B. Create a private Data Exchange so that any teammate can use the worksheet
C. Share the worksheet with teammates within Snowsight
D. Create a database and grant all permissions to teammates

Answer: C

Explanation:
Worksheets in Snowsight can be shared directly with other Snowflake users within the same account. This feature allows for collaboration and sharing of SQL queries or Python code, as well as other data manipulation tasks1.

NEW QUESTION 203
- (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 208
- (Topic 2)
Where can a user find and review the failed logins of a specific user for the past 30 days?

A. The USERS view in ACCOUNT_USAGE
B. The LOGIN_HISTORY view in ACCOUNT_USAGE

C. The ACCESS_HISTORY view in ACCOUNT_USAGE
D. The SESSIONS view in ACCOUNT_USAGE

Answer: B

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

NEW QUESTION 213
- (Topic 2)
What type of query benefits the MOST from search optimization?

A. A query that uses only disjunction (i.e., OR) predicates
B. A query that includes analytical expressions
C. A query that uses equality predicates or predicates that use IN
D. A query that filters on semi-structured data types

Answer: C

Explanation:
Search optimization in Snowflake is designed to improve the performance of queries that are selective and involve point lookup operations using equality and IN predicates. It is particularly beneficial for queries that access columns with a high number of distinct values1.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

NEW QUESTION 214
- (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 215
- (Topic 2)
Which command should be used to download files from a Snowflake stage to a local folder on a client's machine?

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

Answer: B

Explanation:
The GET command is used to download files from a Snowflake stage to a local folder on a client??s machine2.
Reference: https://docs.snowflake.com/en/sql-reference/sql/get.html

NEW QUESTION 219
- (Topic 2)
What is the minimum Snowflake edition required for row level security?

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

Answer: B

Explanation:
Row level security in Snowflake is available starting with the Enterprise edition. This feature allows for the creation of row access policies that can control access to data at the row level within tables and views

NEW QUESTION 221
- (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 223
- (Topic 2)
Which snowflake objects will incur both storage and cloud compute charges? (Select TWO)

A. Materialized view
B. Sequence
C. Secure view
D. Transient table
E. Clustered table

Answer: AD

Explanation:
In Snowflake, both materialized views and transient tables will incur storage charges because they store data. They will also incur compute charges when queries are run against them, as compute resources are used to process the queries. References:
[COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 227
- (Topic 2)
Which minimum Snowflake edition allows for a dedicated metadata store?

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

Answer: B

Explanation:
The Enterprise edition of Snowflake allows for a dedicated metadata store, providing additional features designed for large-scale enterprises
Reference: https://docs.snowflake.com/en/user-guide/intro-editions.html

NEW QUESTION 229
- (Topic 2)
A running virtual warehouse is suspended.
What is the MINIMUM amount of time that the warehouse will incur charges for when it is restarted?

A. 1 second
B. 60 seconds
C. 5 minutes
D. 60 minutes

Answer: B

Explanation:
When a running virtual warehouse in Snowflake is suspended and then restarted, the minimum amount of time it will incur charges for is 60 seconds2.

NEW QUESTION 234
- (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 237
- (Topic 2)
How long is the Fail-safe period for temporary and transient tables?

A. There is no Fail-safe period for these tables.
B. 1 day
C. 7 days
D. 31 days
E. 90 days

Answer: A

Explanation:
Temporary and transient tables in Snowflake do not have a Fail-safe period. Once the session ends or the tables are dropped, the data is purged and not recoverable1.
Reference: https://docs.snowflake.com/en/user-guide/tables-temp-transient.html

NEW QUESTION 239
- (Topic 2)
Which statements are correct concerning the leveraging of third-party data from the Snowflake Data Marketplace? (Choose two.)

A. Data is live, ready-to-query, and can be personalized.
B. Data needs to be loaded into a cloud provider as a consumer account.
C. Data is not available for copying or moving to an individual Snowflake account.
D. Data is available without copying or moving.
E. Data transformations are required when combining Data Marketplace datasets with existing data in Snowflake.

Answer: AD

Explanation:
When leveraging third-party data from the Snowflake Data Marketplace, the data is live, ready-to-query, and can be personalized. Additionally, the data is available without the need for copying or moving it to an individual Snowflake account, allowing for seamless integration with existing data

NEW QUESTION 244
- (Topic 2)
Which of the following objects can be directly restored using the UNDROP command? (Choose two.)

A. Schema
B. View
C. Internal stage
D. Table
E. User
F. Role

Answer: BD

Explanation:
The UNDROP command in Snowflake can be used to directly restore Views and Tables. These objects, when dropped, are moved to a ??Recycle Bin?? where they can be restored within a time limit before they are permanently deleted. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 245
- (Topic 2)
Files have been uploaded to a Snowflake internal stage. The files now need to be deleted. Which SQL command should be used to delete the files?

A. PURGE
B. MODIFY
C. REMOVE
D. DELETE

Answer: C

Explanation:
The SQL command used to delete files from a Snowflake internal stage is REMOVE. This command can be used to remove files from either an internal or external stage within Snowflake

NEW QUESTION 248
- (Topic 2)
Which services does the Snowflake Cloud Services layer manage? (Choose two.)

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

Answer: CE

Explanation:
The Snowflake Cloud Services layer manages various services, including authentication and metadata management. This layer ties together all the different components of Snowflake to process user requests, manage sessions, and control access3.

NEW QUESTION 252
- (Topic 2)
What is the minimum Snowflake edition required to use Dynamic Data Masking?

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

Answer: B

Explanation:
The minimum Snowflake edition required to use Dynamic Data Masking is the Enterprise edition. This feature is not available in the Standard edition2.

NEW QUESTION 254
- (Topic 2)
What do the terms scale up and scale out refer to in Snowflake? (Choose two.)

A. Scaling out adds clusters of the same size to a virtual warehouse to handle more concurrent queries.
B. Scaling out adds clusters of varying sizes to a virtual warehouse.
C. Scaling out adds additional database servers to an existing running cluster to handle more concurrent queries.
D. Snowflake recommends using both scaling up and scaling out to handle more concurrent queries.
E. Scaling up resizes a virtual warehouse so it can handle more complex workloads.
F. Scaling up adds additional database servers to an existing running cluster to handle larger workloads.

Answer: AE

Explanation:
Scaling out in Snowflake involves adding clusters of the same size to a virtual warehouse, which allows for handling more concurrent queries without affecting the performance of individual queries. Scaling up refers to resizing a virtual warehouse to increase its compute resources, enabling it to handle more complex workloads and larger queries more efficiently.

NEW QUESTION 256
- (Topic 2)
Which Snowflake function will interpret an input string as a JSON document, and produce a VARIANT value?

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

Answer: A

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

NEW QUESTION 258
- (Topic 3)
What statistical information in a Query Profile indicates that the query is too large to fit in memory? (Select TWO).

A. Bytes spilled to local cache.
B. Bytes spilled to local storage.
C. Bytes spilled to remote cache.
D. Bytes spilled to remote storage.
E. Bytes spilled to remote metastore.

Answer: AB

Explanation:
In a Query Profile, the statistical information that indicates a query is too large to fit in memory includes bytes spilled to local cache and bytes spilled to local storage. These metrics suggest thatthe working data set of the query exceeded the memory available on the warehouse nodes, causing intermediate results to be written to disk

NEW QUESTION 263
- (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 264
- (Topic 3)
What is used to diagnose and troubleshoot network connections to Snowflake?

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

Answer: A

Explanation:
SnowCD (Snowflake Connectivity Diagnostic Tool) is used to diagnose and troubleshoot network connections to Snowflake. It runs a series of connection checks to evaluate the network connection to Snowflake

NEW QUESTION 267
- (Topic 3)
What internal stages are available in Snowflake? (Choose three.)

A. Schema stage
B. Named stage
C. User stage
D. Stream stage
E. Table stage
F. Database stage

Answer: BCE

Explanation:
Snowflake supports three types of internal stages: Named, User, and Table stages. These stages are used for staging data files to be loaded into Snowflake tables. Schema, Stream, and Database stages are not supported as internal stages in Snowflake. References: Snowflake Documentation1.

NEW QUESTION 268
- (Topic 3)
Which TABLE function helps to convert semi-structured data to a relational representation?

A. CHECK_JSON
B. TO_JSON
C. FLATTEN
D. PARSE_JSON

Answer: C

Explanation:
The FLATTEN table function in Snowflake is used to convert semi- structured data, such as JSON or XML, into a relational format. It expands nested arrays or objects into multiple rows, making the data suitable for relational querying3.

NEW QUESTION 269
- (Topic 3)
How does Snowflake allow a data provider with an Azure account in central Canada to share data with a data consumer on AWS in Australia?

A. The data provider in Azure Central Canada can create a direct share to AWS Asia Pacific, if they are both in the same organization.
B. The data consumer and data provider can form a Data Exchange within the same organization to create a share from Azure Central Canada to AWS Asia Pacific.
C. The data provider uses the GET DATA workflow in the Snowflake Data Marketplace to create a share between Azure Central Canada and AWS Asia Pacific.
D. The data provider must replicate the database to a secondary account in AWS Asia Pacific within the same organization then create a share to the data consumer's account.

Answer: D

Explanation:
Snowflake allows data providers to share data with consumers across different cloud platforms and regions through database replication. The data provider must replicate the database to a secondary account in the target region or cloud platform within the same organization, and then create a share to the data consumer??s account. This process ensures that the data is available in the consumer??sregion and on their cloud platform, facilitating seamless data sharing. References: Sharing data securely across regions and cloud platforms | Snowflake Documentation

NEW QUESTION 273
- (Topic 3)
In which Snowflake layer does Snowflake reorganize data into its internal optimized, compressed, columnar format?

A. Cloud Services
B. Database Storage
C. Query Processing
D. Metadata Management

Answer: B

Explanation:
Snowflake reorganizes data into its internal optimized, compressed, columnar format in the Database Storage layer. This process is part of how Snowflake manages data storage, ensuring efficient data retrieval and query performance

NEW QUESTION 277
- (Topic 3)
What column type does a Kafka connector store formatted information in a single column?

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

Answer: D

Explanation:
The Kafka connector stores formatted information in a single column of type VARIANT. This column type is used to store semi-structured data like JSON or Avro, which allows for flexibility in the data structure

NEW QUESTION 280
- (Topic 3)
Which Snowflake URL type allows users or applications to download or access files directly from Snowflake stage without authentication?

A. Directory
B. File
C. Pre-signed
D. Scoped

Answer: C

Explanation:
The pre-signed URL type allows users or applications to download or access files directly from a Snowflake stage without authentication. This URL type is open and can be used without needing to authenticate into Snowflake or pass an authorization token.

NEW QUESTION 285
- (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 287
- (Topic 3)
Which of the following is the Snowflake Account_Usage.Metering_History view used for?

A. Gathering the hourly credit usage for an account
B. Compiling an account's average cloud services cost over the previous month
C. Summarizing the throughput of Snowpipe costs for an account
D. Calculating the funds left on an account's contract

Answer: A

Explanation:
The Snowflake Account_Usage.Metering_History view is used to gather the hourly credit usage for an account. This view provides details on the credits consumed by various services within Snowflake for the last 365 days1.

NEW QUESTION 288
- (Topic 3)
For the ALLOWED VALUES tag property, what is the MAXIMUM number of possible string values for a single tag?

A. 10
B. 50
C. 64
D. 256

Answer: D

Explanation:
For the ALLOWED VALUES tag property, the maximum number of possible string values for a single tag is 256. This allows for a wide range of values to be

assigned to a tag when it is set on an object

NEW QUESTION 293
- (Topic 3)
Which of the following practices are recommended when creating a user in Snowflake? (Choose two.)

A. Configure the user to be initially disabled.
B. Force an immediate password change.
C. Set a default role for the user.
D. Set the number of minutes to unlock to 15 minutes.
E. Set the user's access to expire within a specified timeframe.

Answer: BC

NEW QUESTION 296
- (Topic 3)
What is the purpose of using the OBJECT_CONSTRUCT function with me COPY INTO command?

A. Reorder the rows in a relational table and then unload the rows into a file
B. Convert the rows in a relational table lo a single VARIANT column and then unload the rows into a file.
C. Reorder the data columns according to a target table definition and then unload the rows into the table.
D. Convert the rows in a source file to a single variant column and then load the rows from the file to a variant table.

Answer: B

Explanation:
The OBJECT_CONSTRUCT function is used with the COPY INTO command to convert the rows in a relational table to a single VARIANT column, which can then be unloaded into a file. This is useful for transforming table data into a semi-structured JSON format

NEW QUESTION 298
- (Topic 3)
What is the MAXIMUM size limit for a record of a VARIANT data type?

A. 8MB
B. 16MB
C. 32MB
D. 128MB

Answer: B

Explanation:
The maximum size limit for a record of a VARIANT data type in Snowflake is 16MB. This allows for storing semi-structured data types like JSON, Avro, ORC, Parquet, or XML within a single VARIANT column. References: Based on general database knowledge as of 2021.

NEW QUESTION 303
- (Topic 3)
Which statement MOST accurately describes clustering in Snowflake?

A. The database ACCOUNTADMIN must define the clustering methodology for each Snowflake table.
B. Clustering is the way data is grouped together and stored within Snowflake micro- partitions.
C. The clustering key must be included in the COPY command when loading data into Snowflake.
D. Clustering can be disabled within a Snowflake account.

Answer: B

Explanation:
Clustering in Snowflake refers to the organization of data within micro- partitions, which are contiguous units of storage within Snowflake tables. Clustering keys can be defined to co-locate similar rows in the same micro-partitions, improving scan efficiency and query performance12.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 304
- (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 305
- (Topic 3)
Which objects together comprise a namespace in Snowflake? (Select TWO).

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

Answer: BC

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

NEW QUESTION 308
- (Topic 3)
What is the recommended way to change the existing file format type in my format from CSV to JSON?

A. READ
B. OWNERSHIP
C. USAGK
D. WRTTF

Answer: A

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

NEW QUESTION 313
- (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 318
- (Topic 3)
What is the recommended way to change the existing file format type in my format from CSV to JSON?

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

Answer: A

Explanation:
To change the existing file format type from CSV to JSON, the recommended way is to use the ALTER FILE FORMAT command with the SET TYPE=JSON clause. This alters the file format specification to use JSON instead of CSV. References: Based on my internal knowledge as of 2021.

NEW QUESTION 322
- (Topic 3)
Which file format will keep floating-point numbers from being truncated when data is unloaded?

A. CSV
B. JSON
C. ORC
D. Parquet

Answer: D

Explanation:
The Parquet file format is known for preserving the precision of floating-point numbers when data is unloaded, preventing truncation of the values3.

NEW QUESTION 323

- (Topic 3)
If a virtual warehouse runs for 61 seconds, shuts down, and then restarts and runs for 30 seconds, for how many seconds is it billed?

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

Answer: D

Explanation:
Snowflake??s billing for virtual warehouses is per-second, with a minimum of 60 seconds for each time the warehouse is started or resumed. Therefore, if a warehouse runs for 61 seconds, it is billed for 61 seconds. If it is then shut down and restarted, running for an additional 30 seconds, it is billed for another 60 seconds (the minimum charge for a restart), totaling 121 seconds2

NEW QUESTION 325
- (Topic 3)
Which Snowflake feature will allow small volumes of data to continuously load into Snowflake and will incrementally make the data available for analysis?

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

Answer: B

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

NEW QUESTION 327
- (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 332
- (Topic 3)
What are advantages clones have over tables created with CREATE TABLE AS SELECT statement? (Choose two.)

A. The clone always stays in sync with the original table.
B. The clone has better query performance.
C. The clone is created almost instantly.
D. The clone will have time travel history from the original table.
E. The clone saves space by not duplicating storage.

Answer: CE

Explanation:
Clones in Snowflake have the advantage of being created almost instantly and saving space by not duplicating storage. This is due to Snowflake??s zero-copy cloning feature, which allows for the creation of object clones without the additional storage costs typically associated with data duplication23. Clones are independent of the original table and do not stay in sync with it, nor do they inherently have better query performance. However, they do inherit the time travel history from the original table at the time of cloning

NEW QUESTION 334
- (Topic 3)
Which privilege is required for a role to be able to resume a suspended warehouse if auto- resume is not enabled?

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

Answer: B

Explanation:
The OPERATE privilege is required for a role to resume a suspended warehouse if auto-resume is not enabled. This privilege allows the role to start, stop, suspend, or resume a virtual warehouse3.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 337
- (Topic 3)
By definition, a secure view is exposed only to users with what privilege?

A. IMPORT SHARE
B. OWNERSHIP
C. REFERENCES
D. USAGE

Answer: B

Explanation:
A secure view in Snowflake is exposed only to users with the OWNERSHIP privilege. This privilege ensures that only authorized users who own the view, or roles that include ownership, can access the secure view

NEW QUESTION 339
- (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 341
- (Topic 3)
What file formats does Snowflake support for loading semi-structured data? (Choose three.)

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

Answer: BDE

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

NEW QUESTION 345
- (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 350
- (Topic 3)
What can a Snowflake user do with the information included in the details section of a Query Profile?

A. Determine the total duration of the query.
B. Determine the role of the user who ran the query.
C. Determine the source system that the queried table is from.
D. Determine if the query was on structured or semi-structured data.

Answer: A

Explanation:

The details section of a Query Profile in Snowflake provides users with various statistics and information about the execution of a query. One of the key pieces of information that can be determined from this section is the total duration of the query, which helps in understanding the performance and identifying potential bottlenecks. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 355
- (Topic 3)
Data storage for individual tables can be monitored using which commands and/or objects? (Choose two.)

A. SHOW STORAGE BY TABLE;
B. SHOW TABLES;
C. Information Schema -> TABLE_HISTORY
D. Information Schema -> TABLE_FUNCTION
E. Information Schema -> TABLE_STORAGE_METRICS

Answer: AE

Explanation:
To monitor data storage for individual tables, the commands and objects that can be used are ??SHOW STORAGE BY TABLE;?? and the Information Schema view ??TABLE_STORAGE_METRICS??. These tools provide detailed information about the storage utilization for tables. References: Snowflake Documentation

NEW QUESTION 356
- (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 361
- (Topic 3)
Which feature allows a user the ability to control the organization of data in a micro- partition?

A. Range Partitioning
B. Search Optimization Service
C. Automatic Clustering
D. Horizontal Partitioning

Answer: C

Explanation:
Automatic Clustering is a feature that allows users to control the organization of data within micro-partitions in Snowflake. By defining clustering keys, Snowflake can automatically reorganize the data in micro-partitions to optimize query performance1.

NEW QUESTION 362
- (Topic 3)
Network policies can be applied to which of the following Snowflake objects? (Choose two.)

A. Roles
B. Databases
C. Warehouses
D. Users
E. Accounts

Answer: DE

Explanation:
Network policies in Snowflake can be applied to users and accounts. These policies control inbound access to the Snowflake service and internal stages, allowing or denying access based on the originating network identifiers12.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 363
- (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 365
- (Topic 3)
How would a user execute a series of SQL statements using a task?

A. Include the SQL statements in the body of the task CREATE TASK mytask .. AS INSERT INTO target1 SELECT .. FROM stream_s1 WHERE .. INSERT INTO target2 SELECT .. FROM stream_s1WHERE ..
B. A stored procedure can have only one DML statement per stored procedure invocation and therefore the user should sequence stored procedure calls in the taskdefinition CREATE TASK mytask .... AScall stored_proc1(); call stored_proc2();
C. Use a stored procedure executing multiple SQL statements and invoke the stored procedure from the tas
D. CREATE TASK mytask .... AS call stored_proc_multiple_statements_inside();
E. Create a task for each SQL statement (e.
F. resulting in task1, task2, etc.) and string the series of SQL statements by having a control task calling task1, task2, et
G. sequentially.

Answer: C

Explanation:
To execute a series of SQL statements using a task, a user would use a stored procedure that contains multiple SQL statements and invoke this stored procedure from the task. References: Snowflake Documentation2.

NEW QUESTION 366
- (Topic 3)
How can a data provider ensure that a data consumer is going to have access to the required objects?

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

Answer: A

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

NEW QUESTION 369
- (Topic 3)
What privilege should a user be granted to change permissions for new objects in a managed access schema?

A. Grant the OWNERSHIP privilege on the schema.
B. Grant the OWNERSHIP privilege on the database.
C. Grant the MANAGE GRANTS global privilege.
D. Grant ALL privileges on the schema.

Answer: C

Explanation:
To change permissions for new objects in a managed access schema, a user should be granted the MANAGE GRANTS global privilege. This privilege allows the user to manage access control through grants on all securable objects within Snowflake2. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 371
- (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 374
- (Topic 3)
When should a user consider disabling auto-suspend for a virtual warehouse? (Select TWO).

A. When users will be using compute at different times throughout a 24/7 period
B. When managing a steady workload

C. When the compute must be available with no delay or lag time
D. When the user does not want to have to manually turn on the warehouse each time it is needed
E. When the warehouse is shared across different teams

Answer: BC

Explanation:
Disabling auto-suspend for a virtual warehouse is recommended when there is a steady workload, which ensures that compute resources are always available. Additionally, it is advisable to disable auto-suspend when immediate availability of compute resources is critical, eliminating any startup delay

NEW QUESTION 379
- (Topic 3)
Which stream type can be used for tracking the records in external tables?

A. Append-only
B. External
C. Insert-only
D. Standard

Answer: B

Explanation:
The stream type that can be used for tracking the records in external tables is ??External??. This type of stream is specifically designed to track changes in external tables

NEW QUESTION 381
- (Topic 3)
For non-materialized views, what column in Information Schema and Account Usage identifies whether a view is secure or not?

A. CHECK_OPTION
B. IS_SECURE
C. IS_UPDATEABLE
D. TABLE_NAME

Answer: B

Explanation:
In the Information Schema and Account Usage, the column that identifies whether a view is secure or not is IS_SECURE2.

NEW QUESTION 382
- (Topic 3)
How do Snowflake data providers share data that resides in different databases?

A. External tables
B. Secure views
C. Materialized views
D. User-Defined Functions (UDFs)

Answer: B

Explanation:
Snowflake data providers can share data residing in different databases through secure views. Secure views allow for the referencing of objects such as schemas, tables, and other views contained in one or more databases, as long as those databases belong to the same account. This enables providers to share data securely and efficiently with consumers. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 384
- (Topic 3)
The first user assigned to a new account, ACCOUNTADMIN, should create at least one additional user with which administrative privilege?

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

Answer: A

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

NEW QUESTION 389
- (Topic 3)
How does a scoped URL expire?

A. When the data cache clears.
B. When the persisted query result period ends.

C. The encoded URL access is permanent.
D. The length of time is specified in the expiration_time argument.

Answer: B

Explanation:
A scoped URL expires when the persisted query result period ends, which is typically after the results cache expires. This is currently set to 24 hours

NEW QUESTION 390
- (Topic 3)
If queries start to queue in a multi-cluster virtual warehouse, an additional compute cluster starts immediately under what setting?

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

Answer: A

Explanation:
In Snowflake, when queries begin to queue in a multi-cluster virtual warehouse, an additional compute cluster starts immediately if the warehouse is set to auto-scale mode. This mode allowsSnowflake to automatically add or resume additional clusters as soon as the workload increases, and similarly, shut down or pause the additional clusters when the load decreases

NEW QUESTION 394
- (Topic 3)
Which features make up Snowflake's column level security? (Select TWO).

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

Answer: BC

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

NEW QUESTION 399
- (Topic 3)
Which task privilege does a Snowflake role need in order to suspend or resume a task?

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

Answer: B

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

NEW QUESTION 404
- (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 406
- (Topic 3)
What does Snowflake's search optimization service support?

A. External tables

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

Answer: C

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

NEW QUESTION 408
- (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 409
- (Topic 3)
What can a Snowflake user do in the Activity section in Snowsight?

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

Answer: D

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

NEW QUESTION 414
- (Topic 3)
What action can a user take to address query concurrency issues?

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

Answer: C

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

NEW QUESTION 417
- (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 421
- (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 424
- (Topic 3)
At what levels can a resource monitor be configured? (Select TWO).

A. Account
B. Database
C. Organization
D. Schema
E. Virtual warehouse

Answer: AE

Explanation:
Resource monitors in Snowflake can be configured at the account and virtual warehouse levels. They are used to track credit usage and control costs associated with running virtual warehouses. When certain thresholds are reached, resource monitors can trigger actions such as sending alerts or suspending warehouses to prevent excessive credit consumption. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 427
- (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 429
- (Topic 3)
Two users share a virtual warehouse named wh dev 01. When one of the users loads data, the other one experiences performance issues while querying data. How does Snowflake recommend resolving this issue?

A. Scale up the existing warehouse.
B. Create separate warehouses for each user.
C. Create separate warehouses for each workload.
D. Stop loading and querying data at the same time.

Answer: C

Explanation:
Snowflake recommends creating separate warehouses for each workload to resolve performance issues caused by shared virtual warehouses. This ensures that the resources are not being overutilized by one user??s activities, thereby affecting the performance of another user??s activities4.

NEW QUESTION 434
- (Topic 3)
Which Snowflake object helps evaluate virtual warehouse performance impacted by query queuing?

A. Resource monitor
B. Account_usag
C. query_history
D. Information_schema.warehouse_load_history
E. Information schema.warehouse metering history

Answer: C

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

NEW QUESTION 435
- (Topic 3)
What role is required to use Partner Connect?

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

Answer: A

Explanation:
To use Partner Connect, the ACCOUNTADMIN role is required. Partner Connect allows account administrators to easily create trial accounts with selected Snowflake business partners and integrate these accounts with Snowflake

NEW QUESTION 438
- (Topic 3)
Where is Snowflake metadata stored?

A. Within the data files
B. In the virtual warehouse layer
C. In the cloud services layer
D. In the remote storage layer

Answer: C

Explanation:
Snowflake??s architecture is divided into three layers: database storage, query processing, and cloud services. The metadata, which includes information about the structure of the data, the SQL operations performed, and the service-level policies, is stored in the cloud services layer. This layer acts as the brain of the Snowflake environment, managing metadata, query optimization, and transaction coordination.

NEW QUESTION 443
- (Topic 3)
Which pages are included in the Activity area of Snowsight? (Select TWO).

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

Answer: DE

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

NEW QUESTION 444
- (Topic 3)
Which statements reflect key functionalities of a Snowflake Data Exchange? (Choose two.)

A. If an account is enrolled with a Data Exchange, it will lose its access to the Snowflake Marketplace.
B. A Data Exchange allows groups of accounts to share data privately among the accounts.
C. A Data Exchange allows accounts to share data with third, non-Snowflake parties.
D. Data Exchange functionality is available by default in accounts using the Enterprise edition or higher.
E. The sharing of data in a Data Exchange is bidirectiona
F. An account can be a provider for some datasets and a consumer for others.

Answer: BE

Explanation:
A Snowflake Data Exchange allows groups of accounts to share data privately among the accounts (B), and it supports bidirectional sharing, meaning an account can be both a provider and a consumer of data (E). This facilitates secure and governed data collaboration within a selected group3.

NEW QUESTION 448
- (Topic 3)
A tabular User-Defined Function (UDF) is defined by specifying a return clause that contains which keyword?

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

Answer: B

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

NEW QUESTION 449
- (Topic 3)

What is the difference between a stored procedure and a User-Defined Function (UDF)?

A. Stored procedures can execute database operations while UDFs cannot.
B. Returning a value is required in a stored procedure while returning values in a UDF is optional.
C. Values returned by a stored procedure can be used directly in a SQL statement while the values returned by a UDF cannot.
D. Multiple stored procedures can be called as part of a single executable statement while a single SQL statement can only call one UDF at a time.

Answer: A

Explanation:
Stored procedures in Snowflake can perform a variety of database operations, including DDL and DML, whereas UDFs are designed to return values and cannot execute database operations1.

NEW QUESTION 454
- (Topic 3)
What type of columns does Snowflake recommend to be used as clustering keys? (Select TWO).

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

Answer: CD

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

NEW QUESTION 456
- (Topic 3)
What is the name of the SnowSQLfile that can store connection information?

A. history
B. config
C. snowsqLcnf
D. snowsql.pubkey

Answer: B

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

NEW QUESTION 460
- (Topic 3)
Which commands should be used to grant the privilege allowing a role to select data from all current tables and any tables that will be created later in a schema? (Choose two.)

A. grant USAGE on all tables in schema DB1.SCHEMA to role MYROLE;
B. grant USAGE on future tables in schema DB1.SCHEMA to role MYROLE;
C. grant SELECT on all tables in schema DB1.SCHEMA to role MYROLE;
D. grant SELECT on future tables in schema DB1.SCHEMA to role MYROLE;
E. grant SELECT on all tables in database DB1 to role MYROLE;
F. grant SELECT on future tables in database DB1 to role MYROLE;

Answer: CD

Explanation:
To grant a role the privilege to select data from all current and future tables in a schema, two separate commands are needed. The first command grants the SELECT privilege on all existing tables within the schema, and the second command grants the SELECT privilege on all tables that will be created in the future within the same schema.

NEW QUESTION 463
- (Topic 3)
A user has a standard multi-cluster warehouse auto-scaling policy in place.
Which condition will trigger a cluster to shut-down?

A. When after 2-3 consecutive checks the system determines that the load on the most- loaded cluster could be redistributed.
B. When after 5-6 consecutive checks the system determines that the load on the most- loaded cluster could be redistributed.
C. When after 5-6 consecutive checks the system determines that the load on the least- loaded cluster could be redistributed.
D. When after 2-3 consecutive checks the system determines that the load on the least- loaded cluster could be redistributed.

Answer: D

Explanation:
In a standard multi-cluster warehouse with auto-scaling, a cluster will shut down when, after 2-3 consecutive checks, the system determines that the load on the

least-loaded cluster could be redistributed to other clusters. This ensures efficient resource utilization and cost management. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 468
- (Topic 3)
Which of the following activities consume virtual warehouse credits in the Snowflake environment? (Choose two.)

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

Answer: BD

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

NEW QUESTION 471
- (Topic 3)
How often are the Account and Table master keys automatically rotated by Snowflake?

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

Answer: A

Explanation:
Snowflake automatically rotates the Account and Table master keys when they are more than 30 days old. Active keys are retired, and new keys are created, ensuring robust security through frequent key changes1

NEW QUESTION 476
- (Topic 3)
Which feature is integrated to support Multi-Factor Authentication (MFA) at Snowflake?

A. Authy
B. Duo Security
C. One Login
D. RSA SecurID Access

Answer: B

Explanation:
Snowflake integrates Duo Security to support Multi-Factor Authentication (MFA). This feature provides increased login security for users connecting to Snowflake, and it is managed completely by Snowflake without the need for users to sign up separately with Duo4.

NEW QUESTION 478
.....

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